A PRECISE SEMANTICS FOR VAGUE DIAGRAMS

TIM MENZIES,
PAUL COMPTON
AI Lab, School of Computer Science And Engineering,
University of NSW, P.O. Box 1, Kensington, NSW, Australia, 2033
Email: {timm|compton}@cse.unsw.edu.au; fax: 02-663-4576

ABSTRACT

Informal vague causal diagrams (VCDs) are a common technique for illustrating and sharing expert intuitions. Normally, VCDs are viewed as precursors to other modelling techniques which necessitates further knowledge acquisition. Here we explore what semantics can be granted to VCDs, without having to request more information from the expert(s) or the domain. The impreciseness of VCDs typically makes them indeterminate. VCD inferencing must assume multiple possibilities and manage mutually exclusive possibilities in separate worlds. Given a library of known behaviour of the entity being modelled, we can use *exhaustive abduction* over VCDs to prove what behaviours are categorically impossible; i.e. we can use VCDs for knowledge acquisition. *KEYWORDS:* Common sense reasoning, knowledge acquisition, knowledge representation, knowledge sharing technology, qualitative reasoning, reasoning about physical systems, situated cognition, truth maintenance, diagrammatic reasoning, causality, abduction

1. Introduction

Diagrammatic reasoning (DR) is a poorly defined field. Despite attempts to define general principles for DR[1], any summary of recent work in the area[2] demonstrates that (i) DR means very different things to different researchers; and (ii) that the field lacks unifying theories or principles. However, we can approximately divide the field into two camps: propositional/sentential vs image/analogical (though some overlap exists[3]). The slogan of analogue DR research is "people don't do inference, the world does most of it for us". Analogue DR[4] uses properties of the representation to avoid inferencing. For example, to work out what cities lie under a great-circle route between London and Sydney, take a globe and run a finger over the desired route, writing down what cities you touch[5]. Sentential DR typically converts a two-dimensional representation into a sentential form that can be processed by some form of logic. Significantly for the sentential vs analogue debate, logical proofs can be generated faster using heuristics taken from visual inferencing[6]. Further, convincing and intuitive proofs of certain theorems are trivial and rapid to produce visually but require non-trivial sentential inference[7]. Claims that humans use diagrams internally to augment/ replace logical inferencing are common[8] but this is not universally accepted[9].

Despite the shortcuts offered by the analogue approach, we endorse sentential DR. Methodologicaly, we believe it more useful to explore general principles rather than specific devices for special domains. Although there exist examples of general principles for analogue DR[10], these are few in number. The sentential framework presented here makes sense of diagrams using a general inference procedure with wide applicability (exhaustive abduction).

The kind of diagram we will explore are the vague causal diagrams (VCDs) drawn to share expert intuitions in domains like neuroendocrinology (the study of the interaction of nerves and glands). Such diagrams consist of nodes connected by arcs labelled (e.g.) "inhibits", "+", "promotes", "-", or "blocks". These diagrams are common (e.g. figure 1). Our neuroendocrinological expert could find five such graphs in as many minutes from the first two textbooks he took randomly from his bookshelf.

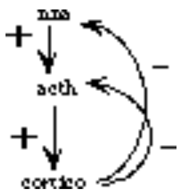

Figure 1: *VCD for connections between serum adrenocorticotropin (*acth*), serum corticosterone (*cortico*), and neuro-noradrenergic activity (*nna *- measured as the ratio of noradrenaline to its post-cursor, 3,4-dihydroxphenylethethyleneglycol). VCD drawn by Smythe*[11].

In this paper, we offer a novel sentential definition for diagramatic reasoning based on our research on knowledge acquisition for neuroendocrinology. We say that such a VCD is *understood* iff we can extract from it a deductive theory that can explain some of our known behaviour without also entailing inconsistencies. We understand that VCD_x is better than VCD_y iff the theory extracted from VCD_x can explain more known behaviours than the theory extracted from VCD_y . This extraction process is defined in §3, after a general discussion in §2 about VCDs. §4 discusses possible limits with our sentential diagram understanding system. §5 describes our experimental results (e.g. sentential VCD understanding can yield insights that are invisible to other techniques). §6 discusses related work.

Applications for this novel definition include multiple-expert knowledge acquisition (when feuding experts need some judgment about competing knowledge), group decision support systems (when groups use vague diagrams as a knowledge sharing tool), and single-expert knowledge acquisition (when a single expert is unsure about what knowledge to add to a knowledge base). Note that this work is a generalisation of Compton & Feldman's qualitative hypothesis testing project[12,13].

2. About Vague Causal Diagrams (VCDs)

VCDs are usually viewed as pre-cursors to a more formal modelling technique. For example, vague statements such as "glucose levels effect insulin production" can be translated into a *compartmental model:* a set of exponential functions that control flows in and out of the insulin "compartment". This translation typically requires more information that what is available in the original diagram (e.g. numeric parameters for the exponential functions modelling the flow rates). In poorly measured domains (e.g. economics, ecology, and most of human internal medicine including neuroendocrinology) this information may be currently unavailable (e.g. in the case of VCDs drawn as hypothesis about new ideas) or prohibitively expensive to collect.

These limitations with numeric modelling have lead some mathematical modellers to question using quantitative methods for understanding the kinds of diagrams we call VCDs. For example, after producing eight different quantitative models for the same phenomena (human ovulation), McIntosh & McIntosh[14] comment that:

> *The most striking feature evident from studying these models is the variety of equations which give reasonable representations of the observed experimental data... In each model these apparently appropriate equations have been derived from quite different assumptions and simplifications and use different parameters.*

McIntosh & McIntosh believe that deeper than the non-unique empirical models is another kind of model which embodies our actual concept of what caused the observed behaviour. We concur and argue that, wherever possible, the natural expression of the domain (i.e. VCDs) should not be inappropriately contorted into a numeric formalism.

Processing VCDs directly without further knowledge acquisition is complicated by qualitative indeterminacy. For example, in figure 1, when both *cortico* and *nna* go ↑ (up), we have two competing influences on *acth* , the net effect of which could be *acth*↑ , *acth*↓ (down), or *acth*θ (remains steady). This indeterminacy is worse in poorly measured domains since we can't tame indeterminacy by using known measurements to reject possibilities. Also, when generating explanations, we will often make assumptions about unmeasured vertices. Some of these assumptions will be incompatible. VCD inference must manage the incompatible assumptions in different *worlds* (defined below). Lastly, given the informal manner of VCD construction, only a subset of a VCD (*VCD'*) may be consistent and *VCD'* may only be able to explain a subset of known behaviour. VCD inference must be a search for some subset of the total diagram that can explain some subset of known effects without generating inconsistencies (e.g. proving *acth*↑ and *acth*↓ simultaneously). Given the indeterminate nature of the inference, there may be multiple subsets and we will have to define some criteria for choosing between them. Formally, VCD inference is a variant on abduction (see next section).

3. About Abduction

Consider a system with two facts *a*, *b* and a rule R_1: *If* $a \Rightarrow b$. *Deduction* is the inference from *a* to *b*. *Induction* is the process of learning R_1 given examples of *a* and *b* occurring together. *Abduction* is inferring *a*, given *b*[15]. Abduction is only a plausible inference since other rules may have concluded *b* using other premise(s). Hence abduction requires some inference assessment operator (which we will call *BEST*). Abduction is a mechanical procedure that, at a symbol level, is equivalent for a variety of "⇒" operators: logical deduction, causality, or application of a default rule[16] (however, various authors caution against mixing up "⇒" operators within a single knowledge base[17,18]). Intuitively, abduction is the generation and evaluation of possible behaviours[19] and has been applied to many domains; e.g. diagnosis[20,21], causal reasoning[22], natural language processing[23], explanation generation[24], planning, and design[25], visual pattern recognition[26], frame-based reasoning[27,28] and case-based reasoning[29].

The abductive problem we will consider is a tuple <*M', I, IN, OUT, BEST*> defined as follows. Let *M* be some model and *M'* be a directed, possibly cyclic graph generated via converting *M* into a propositional form (for example, see Figure 2). *M'* is the space of possible ground proof trees that could be extracted from *M*. More precisely: (i) the vertices *V* of *M'* are either literals from *M* or *Ands* (defined below); (ii) the directed edge $E_{xy}(From_x, To_y)$ of *M'* represents an inference rule of the form: $E_w(V_x,V_y)$ iff literal V_x could explain literal V_y.

We translate *M* to *M'* since (i) the translation converts any domain-specific processing to (possibly) additional edges and/or vertices; hence, (ii) lets us define our adductive process in terms of a uniform *M'* structure; (iii) we can define macro-expansions for commonly used sub-graphs; (iv) we can assign each *M'* vertex a unique integer id and use bitstrings to optimise the inferencing. Note that *M'* explicitly represents the search space tacit in *M*.

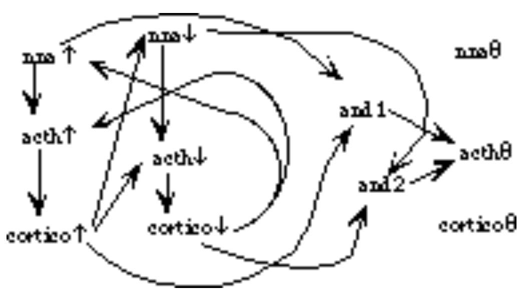

Figure 2: M' *for the* M *of figure 1. One* M' *vertex has been created for each possible state of* M *vertices (*up, down, steady*).* Ands *have been added to combine influences that can lead to a steady (e.g.* nna↑ & cortico↑ ⇒ acthθ*). The vertices* nnaθ *and* corticoθ *are isolated since no combination of influences can combine to lead to these steady vertices.*

Let *IN* and *OUT* be subsets of the *M'* vertices. *P* is the set of non-cyclic proofs P_X whose nodes, edges, leaves, roots are from *V*, *E, IN, OUT* respectively and whose edges share a vertex with at least one other edge in P_X. No two vertices in P_X can violate the invariants *I*. For all *Ands* in P_X, all the in-edges of that *And* must also be in *P* ; i.e. all parents of that *And* are also in P_x. For all *V* that are not *Ands*, P_X must contain zero or one parent only. *OUT* and *IN* list the known measurements of the entity being modelled. Vertices in *P* that are not from *OUT* or *IN* are assumptions *A*.

Abduction is the generation of a world W_X; i.e. the edges contained in a union of a subset of *P* such that (i) $DONE \subset OUT$; (ii) $USED \subset IN$; (iii) W_X & $USED \Rightarrow DONE$; (iv) $\neg(W_X$ & $USED \Rightarrow$ false) i.e. does not violate *I*. ; (v) inference for (ii) and (iii) is restricted to W_X; (vi) W_X is maximal with respect to set inclusion of *E*. For example, consider an abduction over the *M'* of figure 2 with *IN* = {*cortico*↑}, *OUT* = {*nna*↓, *acth*θ}, and *I* being the rule that P_X cannot include an

↑, ↓, or θ simultaneously for the same M vertex. We would generate one world $W_1 = \{P_1\} = \{E_1(cortico\uparrow, nna\downarrow)\}$ with $A = \{\}$, $USED = IN$ and $DONE = \{nna\downarrow\}$. $Acth\theta$ is not in $DONE$ since its explanation requires a proof through two different states of either *cortico* or *acth* (which would violate *I*).

Exhaustive abduction (EA) is the generation of all W_X and their subsequent evaluation by a domain-specific *BEST* operator. Example *BESTs* include $BEST_1$: returning all W_X with fewest assumptions ($|A|$); $BEST_2$: with fewest number of causes ($|USED|$); $BEST_3$: with shortest proof size ($|P_x|$); $BEST_4$: with the largest number of explained effects ($|DONE|$); or $BEST_5$ which avoids edges with low likelihood (assuming that such meta-knowledge about edges is available; e.g. some edges were proposed as part of a theory you wish to fault).

Our definition of abduction is compatible with abduction as defined elsewhere,[30,31,32]. Our worlds are different to the *extensions* of Reiter's default logic[33] in that not all consequences of literals in W_X appear in W_X: if a proposition is not in a path to a member of *OUT*, it is ignored (this saves unnecessary world generation). Note that W_X satisfies our requirement for understanding VCDs: (i)W_X is some portion of M which can be used to simply infer *DONE* from *USED*; (ii)W_X contains consistent assumptions; (iii) *BEST* is an explicit selection criteria.

4. Apparent Limitations

This section discusses possible limitations to our approach. In practice, these issues are not overly restrictive (see below).

M must be a finite theory otherwise M' (the explicit search space tacit in M) will be infinite and EA will never terminate. P_x can't contain loops. Hence, we cannot explain time-series data (e.g X went up, then later it went down, then it went up again) without significantly increasing the size of our models (i.e. create one M' vertex for each literal in M at all measured time intervals). All our experiments (see below) assume non-time-series data.

General abduction is known to be NP-hard[34,35] and polynomial time abductive inference procedures are only known for certain restrictive cases: e.g. unit resolution over non-cyclic background theories[36]; or when sufficient "rule-out knowledge" is available to cull much of the search space[37]. EA does more work than standard abduction and so we could pessimistically predict that the generation of all W_X to be impractically slow.

Another pre-experimental pessimistic prediction about EA would be that any behaviour can be generated from a search through indeterminate models. If so, then the power of our sentential DR approach would be minimal since it would incorrectly "understand" that any VCD can do anything.

5. Experimental Work

Compton & Feldman used $BEST_4$ to demonstrate that a VCD for glucose regulation developed from international refereed publications[38] could not produce proofs for a large number of known causes and effects. In all, 109 of 343 (32%) of the known data points from six studied papers could not be explained with reference to their model. Of these detected faults, at least one represented an insight into the process of glucose regulation that had been invisible to conventional scientific review process[39]. Interestingly, the faults were detected using data published to support the models; i.e. it is possible that waiting in all the publications in the books in all the libraries around the world is a backlog of extra inferences that we could make about existing knowledge, without having to perform expensive further experimentation.

A subsequent study by Menzies[40] corrected some modelling errors of Compton & Feldman to increase the inexplicable percentage from 32% to 45%. Another smaller study used $BEST_4$ to successfully fault the published scientific theory of figure 1[41].

Compton & Feldman's system, and two subsequent prototypes, used a basic chronological backtracking approach (i.e. no memoing) that was very slow. Basic chronological backtracking has the disadvantage that any feature of the space learnt by the search algorithm is forgotten when backtracking on failure. The current implementation, HT4, caches the most-upstream contraversial assumptions as a side-efect of proof generation (in a manner analagous to the generation of ATMS minimal labels[42]). This system runs 130 times faster than the Compton & Feldman system since world switching does not require extensive further computation[43]. Menzies & Gambetta reports 2 studies involving 4504 EA runs over 299 different *M'* models[44]. The *Changing N study* artificially generated 94 *M'* of varying numbers of vertices (*N*) while keeping *B* (the average number of children/vertex) constant at 2.25. The *Changing B study*, artificially generated 205 *M'* of varying fanout while keeping *N* constant at 554. Figure 3 shows the results.

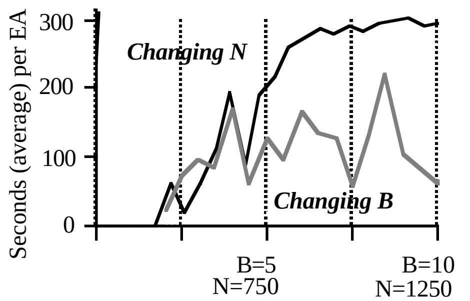

Figure 3A: *Runtimes from Menzies & Gambetta. Note that the plateau after* N*=800 is an artefact of the "give-up" time limit of 5 minutes built into HT4 (300 seconds).*

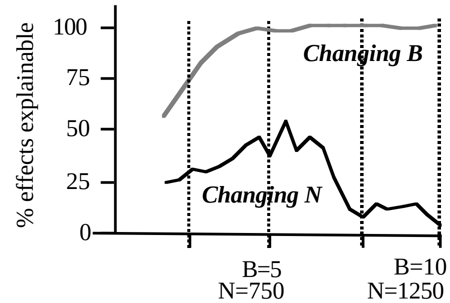

Figure 3B: *% effects explainable from Menzies & Gambetta. Figure 3a shows that after* N*=800, most of the runs did not terminate before the "give-up" time of 300 seconds. Hence, the* changing N *curve in this figure drops off suddenly after* N*=800 .*

Experimentally, we see that HT4 is limited to *M'* of $N < 800$ and $B < 4$. The *N* limit is a function of the current implementation but the *B* limit may be fundamental to exhaustive abduction. Clearly, after a certain level of inter-connectivity, it would be possible to find proofs for any behaviour. Figure 3B suggests that limit is $B = 4$.

Based on known sizes of fielded expert systems[45], Menzies & Gambetta argue that these limits are larger than models we see in contemporary knowledge engineering practice; i.e. we can scale up VCDs to knowledge bases at least as big as those seen in current practice. Further, the level of critique offered by EA can be non-trivial. The changing *N* study of figure 3B shows that up to 75% of behaviour may be falsifiable.

6. Related Work

The internal structures of some validation tools for propositional systems[46,47] use a multiple-worlds architecture which, like ours, were inspired by the simultaneous context generation of DeKleer's ATMS. Our focus is the validation of hastily scribbled diagrams of qualitative domains, though our techniques also apply to propositional systems. These other systems also incorporate model modification operators. We do not explore model modification since we foresee that various users would treasure their favourite portions of their model(s) (typically, the ones they have developed and successfully defended from all critics). It would be unacceptable to permit a learning algorithm scribble all over this knowledge. Learning programs for this domain must strive to preserve the current background theory (an approach explored by Mahidadia in an ILP framework[48,49]).

Qualitative reasoning (QR) processes VCDs by converting them into systems of equations whose numeric values are replaced by one of three qualitative states: *up*, *down*, or *steady*[50]. Early QR was limited to systems that were piece-wise well-approximated by low-order linear equations or by first-order non-linear differential equations (e.g. CONFLUENCES[51], QSIM[52], QPT[53]). Subsequent work studied the asymptotic long-term behaviour of more complex systems of equations[54,55]. Re-expressing

VCDs in qualitative equational terms may not necessarily reflect the intuitions of human modellers[56]. Mapping VCDs to mathematics was required when no alternative execution format was available. Now we offer an alternative approach.

Belief networks (BNs) deduce causality from a statistical analysis of the frequency distributions of variables in a sample to deduce acyclic "networks" (which are really trees) of causal relationships between variables[57]. BNs assume sufficient measurements are available for the statistical analysis; i.e. they are inappropriate in poorly measured domains such as neuroendocrinology. Also, current state-of-the-art BNs assumes acyclic models[58] and models in our test domain are usually cyclic (e.g. figure 1). Further, the theories generated by BNs make do not preserve current beliefs (see above remarks regarding preserving the background theory).

Causality was a central concern in QR in the mid-80s[59]. However, after an inconclusive public debate[60], the term was avoided by experienced researchers (exception: the BN community). Causality in QR was usually a deduced concept from the equations of the qualitative model[61]. This causal deduction process is unsuitable for our domain: since our models may be wrong, the deduced causal graphs may also be wrong. An alternative approach was taken by the "causal editors"; i.e. those researchers like us that let their experts specify causal graphs directly. Early work (e.g. CASNET[62]) demonstrated the utility of this approach. Subsequent work (e.g. ABEL[63]) argued for adding abstraction hierarchies to causal models. Causal models at various levels of abstraction permit inferencing down/up/across abstraction level(s) if more/less/same abstraction is useful in the reasoning. We do not use CASNET-style causation strengths on our edges since like most numbers in our domain, these strengths are unknown. Nor do we use ABEL-style abstraction hierarchies since in our search for all possible answers, we will explore the entire theory, across every abstraction level; i.e. a bigger search. We also address a different problem to other causal editing researchers. The primary use for our models is verification, not simulation (as with CONFLUENCES) or diagnosis (e.g ABEL, CASNET).

7. Conclusion

We have described general principles for sentential DR: (i) permit experts to draw vague causal diagrams; (ii) translate them into propositional theories with a set of invariants; (iii) critique them via libraries of known behaviour (a set of pairs <*IN*, *OUT*>). We say that we understand that diagram iff we can extract from it a deductive theory that can explain some of our known behaviour without also entailing inconsistencies. Further, we understand that some diagram is better other diagram(s) iff its extracted theory explains more known behaviours than its competitor(s). At the symbol-level, this theory extraction process is exhaustive abduction (EA).

Our experimental results demonstrate that we can offer a non-trivial level of critique; i.e. *EA is a useful technique for knowledge acquisition*. The current implementation is limited to critiquing non-cyclic behaviour (e.g. no time-series) from finite models with an average fanout < 4 and $|V|$ in M' < 800. These limits are beyond the models that we can find documented in current knowledge engineering practice; i.e. *EA is a practical technique*. Further, EA over VCDs can identify errors that are invisible to alternative methods; i.e. *EA may be a superior technique of understanding diagrams than alternative quantitative methods*.

Note that we understand our diagrams using general logical principles as apposed to the ad-hocary of standard analogue DR.

8. References

1 Larkin, J., Simon, H.A. *Why a Diagram is (Sometimes) Worth 10,000 Words*, **Cognitive Science**, 1987, 11, p65-99.

2 Narayanan, N.H., **AAAI Spring Symposium on Reasoning with Diagrammatic Representations.** 1992, Stanford University:

3 Myers, K.L. and K. Konolige. *Integrating Analogical and Sentential Reasoning for Perception*

in reference #2.
4 Funt, B.V., *Problem-Solving with Diagrammatic Representations* **Artificial Intelligence**, 1980. **13**: p. 201-230.
5 Funt, B.V. *Experiential Reasoning* in reference #2.
6 Barwise, J. and J. Etchemendy. *Hyperproof: Logical Reasoning with Diagrams* in reference #2.
7 Barker-Plummer, D. and S.C. Bailin. *Proofs and Pictures: Proving the Diamond Lemma with the GROVER Theorem Proving System* in reference #2.
8 Kindfield. *Expert Diagrammatic Reasoning in Biology* in reference #2.
9 Slezak, P. *The "Philosophical" Case Against Visual Images: a "Crucial" Experiment* in reference #2.
10 Cheng, P.C. *Diagrammatic Reasoning in Scientific Discovery: Modelling Galileo's Kinematic Diagrams* in reference #2.
11 Smythe, G.A., Hypothalamic noradrenergic activation of stress-induced Adrenocorticotropin (ACTH) release: Effects of acute and chronic dexamethasone pre-treatment in the rat, **Exp. Clin. Endocrinol. (Life Sci. Adv.)**, 1987. (6): p. 141-144.
12 Feldman, B.T., P.J. Compton, and G.A. Smythe. *Towards Hypothesis Testing: JUSTIN, Prototype System Using Justification in Context* in **Proceedings of the Joint Australian Conference on Artificial Intelligence, AI '89**. 1989.
13 Feldman, B.Z., P.J. Compton, and G.A. Smythe. *Hypothesis Testing: an Appropriate Task for Knowledge-Based Systems* in **4th AAAI-Sponsored Knowledge Acquisition for Knowledge-based Systems Workshop**. 1989. Banff, Canada, October 1989.:
14 McIntosh, J.E.A. and R.P. McIntosh, **Mathematical Modelling and Computers in Endocrinology**. 1980, Berlin: Springer-Verlag.
15 Levesque, H. *A Knowledge-Level Account of Abduction (preliminary version)* in **IJCAI '89**. 1989. Detroit, Michigan, USA:
16 Bylander, T., D. Allemang, M.C. Tanner, and J.R. Josephson, *The complexity of abduction* **Artificial Intelligence**, 1991. **49**: p. 25-60.
17 Charniak, E. and D. McDermott, **Introduction to Artificial Intelligence**. 1987, Addison-Wesley. 701.
18 Poole, D. *Normality and Faults in Logic-Based Diagnosis* in **IJCAI '89**. 1989. Detroit, USA:
19 O'Rourke, P., Working Notes of the 1990 Spring Symposium on Automated Abduction. 1990, University of California, Irvine, CA.:
20 Console, L. and P. Torasso, *A spectrum of definitions of model-based diagnosis* **Computational Intelligence**, 1991. **7**(3): p. 133-141.
21 Pople, H.E. On the Mechanization of Abductive logic, IJCAI '73, p147-152.
22 See Poole, Charniak, Bylander and Levesque references. (#18, #17, #35, #15).
23 See multiple examples in the O'Rourke reference (#19).
24 Leake, D.B., *Goal-Based Explanation Evaluation* **Cognitive Science**, 1991. **15**: p. 509-545.
25 Poole, D. *Hypo-Deductive Reasoning for Abduction, Default Reasoning, and Design* in **Working Notes of the 1990 Spring Symposium on Automated Abduction**. 1990. UC Irvine.
26 Poole, D., *A Methodology for Using a Default and Abductive Reasoning System* **International Journal of Intelligent Systems**, 1990. **5**: p. 521-548.
27 Ramsey, C.L., J.A. Reggia, D.S. Nau, and A. Ferrention, *A Comparative Analysis of Methods for Expert Systems* **Int. J. Man-Machine Studies**, 1986. **24**: p. 475-499.
28 Reggia, J.A. *Abductive Inference* in **Proceedings of the Expert Systems in Government Symposium**. 1985. Washington, D.C.:
29 Leake, D.B. *Focusing Construction and Selection of Abductive Hypotheses* in **IJCAI '93**. 1993.
30 Selman, B. and H.J. Levesque. *Abductive and Default Reasoning: a Computational Core* in **AAAI '90**. 1990.
31 See Poole 1990 (#25, #26).
32 Eshghi, K. *A Tractable Class of Abductive Problems* in **IJCAI '93**. 1993. Chambery,

France:
[33] Reiter, R., *A Logic for Default Reasoning* **Artificial Intelligence**, 1980. **13**: p. 81-132.
[34] See the Selman reference (#30).
[35] Bylander, T., D. Allemang, M.C. Tanner, and J.R. Josephson, *The complexity of abduction* **Artificial Intelligence**, 1991. **49**: p. 25-60.
[36] See the Eshghi reference (#32).
[37] See the Bylander reference (#35)
[38] Smythe, G.A., *Brain-hypothalmus, Pituitary and the Endocrine Pancreas,* in **The Endocrine Pancreas**, S. R., Editor. 1989, Raven Press: New York.
[39] See the Compton & Feldman references (#12, #13).
[40] Menzies, T., **Principles of Generic Testing as Abduction** (in preparation). 1994, PhD thesis, New South Wales University.
[41] Menzies, T., A. Mahidadia, and P. Compton. *Using Causality as a Generic Knowledge Representation, or Why and How Centralised Knowledge Servers Can Use Causality* in **Proceedings of the 7th AAAI-Sponsored Banff Knowledge Acquisition for Knowledge-Based Systems Workshop**. 1992. Banff, Canada, October 11-16,:
[42] DeKleer, J., *An Assumption-Based TMS* **Artificial Intelligence**, 1986. **28**: p. 163-196.
[43] Menzies, T.J. *The Complexity of Model Review* in **Dx -93: The International Workshop on Principles on Model-Based Diagnosis**. 1993. Absersywath, Wales, UK:
[44] Menzies, T.J., Gambetta, W. *Exhaustive Abduction: A Practical Model Validation Tool* in **ECAI '94 Workshop on Validation of Knowledge-Based Systems**. 1994. Amsterdam, Holland:
[45] Preece, A.D. and R. Shinghal. *Verifying Knowledge Bases by Anomaly Detection: An Experience Report* in **ECAI '92**. 1992. Vienna:
[46] Ginsberg, A. *A new approach to checking knowledge bases for inconsistentcy and redundancy* in **Proc. 3rd Annual Expert Systems in Government Conference**. 1987. IEEE Computer Society.
[47] Zlatareva, N. *Distributed Verification and Automated Generation of Test Cases* in **IJCAI '93 workshop on Validation, Verification and Test of KBs**. 1993. Chambery, France:
[48] Mahidadia, A.J., C. Sammut, P. Compton, *Applying Inductive Logic Programming to Causal Qualitative Models in Neuroendocrinology,* **Proceedings of AI '94**. Armidale, NSW, November, 1994.
[49] Mahidadia, A.J., C. Sammut , and P. Compton. *Building and Maintaining Causal Theories* in **AAAI Symposium on Knowledge Assimilation**. 1992. Stanford University, Spring, 1992.:
[50] Iwasaki, Y., *Qualitative Physics,* in **The Handbook of Artificial Intelligence**, A. Barr, P.R. Cohen, and E.A.. Feigenbaum, Editor. 1989, Addison Wesley: p. 323-413.
[51] DeKleer, J. and J.S. Brown, *A Qualitative Physics Based on Confluences* **Artificial Intelligence**, 1984. **25**: p. 7-83.
[52] Kuipers, B., *Qualitative Simulation* **Artificial Intelligence**, 1986. **29**: p. 229-338.
[53] Forbus, K., *Qualitative Process Theory* **Artificial Intelligence**, 1984. **24**: p. 85-168.
[54] Ishida, Y. *Using Global Properties for Qualitative Reasoning: A Qualitative System Theory* in **Proceedings of IJCAI '89**. 1989.
[55] Yip, K.M., *Understanding Complex Dynamics by Visual and Symbolic Reasoning* **Artificial Intelligence**, 1991. **51**: p. 179-221.
[56] Iwasaki, Y. and S. H.A., *Retrospective on "Causality in device behaviour"* **Artificial Intelligence**, 1993. **59**: p. 141-146.
[57] Pearl, J. and T.S. Verma. *A Theory of Inferred Causation* in **Principles of Knowledge Representation and Reasoning, Proceedings of the Second International Conference**. 1991. Morgan Kaufmann.
[58] Geigner, D., A. Paz, and J. Pearl, *Learning Simple Causal Structures* **International Journal of Intelligent Systems**, 1993. **8**: p. 231-247.
[59] Williams, B.C. and J. DeKleer, *Qualitative Reasoning about Physical Systems: a Return to Roots* **Artificial Intelligence**, 1991. 51: p. 1-9.
[60] Iwasaki, Y. and H. Simon, *Theories of Causal Ordering: Reply to deKleer and Brown* **Artificial Intelligence**, 1986. **29**: p. 63-72.

[61] Iwasaki, Y. and H.A. Simon, *Causality in Device Behaviour* **Artificial Intelligence**, 1986. 29: p. 3-31.
[62] Weiss, S.M., C.A. Kulikowski, and S. Amarel, *A Model-Based Method for Computer-Aided Medical Decision-Making* **Artificial Intelligence**, 1978. **11**(145-172).
[63] Patil, R.S., P. Szolovitis, and W.B. Schwartz. Causal Understanding of Patient Illness in Medical Diagnosis in IJCAI '81. 1981.