Diagnostic Improvement Through Qualitative Sensitivity Analysis and Aggregation

Keith L. Downing *
Department of Computer and Information Sciences
The University of Oregon
(503) 686-4408
CSNET: downing@uoregon

Abstract

This paper lays the foundation for a diagnostic system that improves its performance by deriving symptom-fault associations from an underlying causal model and then utilizes those relationships to impose further structure upon the "deep" model. A qualitative version of sensitivity analysis is introduced to extract the implicit symptom-fault information from a set of local constraints. Parameter aggregation triggered by this new information then simplifies diagnosis by forming a more abstract causal representation. The resulting diagnostician thus employs both an experiential and a first-principle approach, where in this case "experiences" are compiled directly from first-principles. Key issues include the roles of knowledge compilation and abstraction in refining qualitative models of physical systems.

Submitted to AAAI 1987 Conference
Science Track
Area: Expert Systems
Subarea: Diagnosis
Word Count: 3300

*This work was partially supported by both a FIPSE grant, number G008440474-02, and a Tektronix Graduate Fellowship.

1 Introduction

Reiter (1987) recognizes two approaches to automated diagnosis:

1. Experiential methods in which direct symptom-fault links distilled from human expert knowledge facilitate quick diagnoses requiring little or no in-depth causal reasoning.

2. First-principle reasoning whereby explict "deep" system models are used to derive the causal pathways from faults to symptoms.

Drawbacks of the experiential method include the generation of multiple fault hypotheses – often resolvable only by weak probabilistic means – and limited explanation capabilities. First-principle diagnosis provides causal explanations at the price of extensive reasoning and/or simulation. However, a deep-model diagnostic system that retains its derived symptom-fault associations can reduce future diagnostic effort without sacrificing explanation abilities.

Our current research involves the compilation of symptom-fault relationships from a mechanical model of the circulatory system with the intention of reusing that information to simplify later diagnosis and therapy. We have partitioned the acquisition of diagnostic skill into two stages:

1. The derivation of symptom-fault connections by applying constraint satisfaction to the causal model.

2. The use of these associations to support the aggregation of structures and parameters to simplify the original model.

Through this process, an automated diagnostician can acquire both rational heuristics supported by an underlying causal model and useful abstractions of that model. This avoids the standard expert-system dependence upon shallow, ad hoc rules and ill-defined symptom and disease hierarchies. This paper discusses a qualitative version of Campbell's (1983) sensitivity analysis as a knowledge-compilation methodology for simplifying qualitative reasoning about complex physical systems.

2 The Cardiovascular Model

The quantitative cardiovascular model developed by Peterson and Campbell(1985) serves as the physical system to undergo diagnosis. In this simulation environment, observable parameters are partitioned into "properties" and "variables". The former represent the relatively static values of a real circulatory system such as vascular resistance to blood flow, or heart strength. Property deviations constitute "faults" and result only from the actions of external factors not represented in the model. Hence, they are always independent parameters in causal relationships. Variables, such as cardiac blood flow or atrial pressure shift value either in direct response to property changes or indirectly through other variable changes. In either case, they represent the dependent system parameters whose deviations constitute "symptoms".

Figure 1 (Peterson and Campbell, 1985) portrays the basic circulatory topology. Briefly, the left heart pumps "clean" oxygen-rich blood through the systemic ("bodily") arteries to the body's capillary beds -- the narrow capillaries being the major source of resistance to blood flow. So blood that exited the left heart at a pressure of approximately 190 mm Hg returns to the right heart via the systemic veins at a pressure close to 5 mm Hg. This "dirty" carbon-dioxide-laden blood gets pumped through the pulmonary arteries to the lungs, where it becomes re-oxygenated before returning to the left heart.

After filling with de-pressurized blood, the left and right hearts, playing the roles of potential generators, contract with both inflow and outflow valves closed, thereby building enough pressure to open the arterial valve and send blood flowing toward the body and lungs respectively. Although the blood volume of the circulatory system remains constant except during bleeding, the amount of "active" blood in the systemic

Figure 1: Basic Circulatory Topology

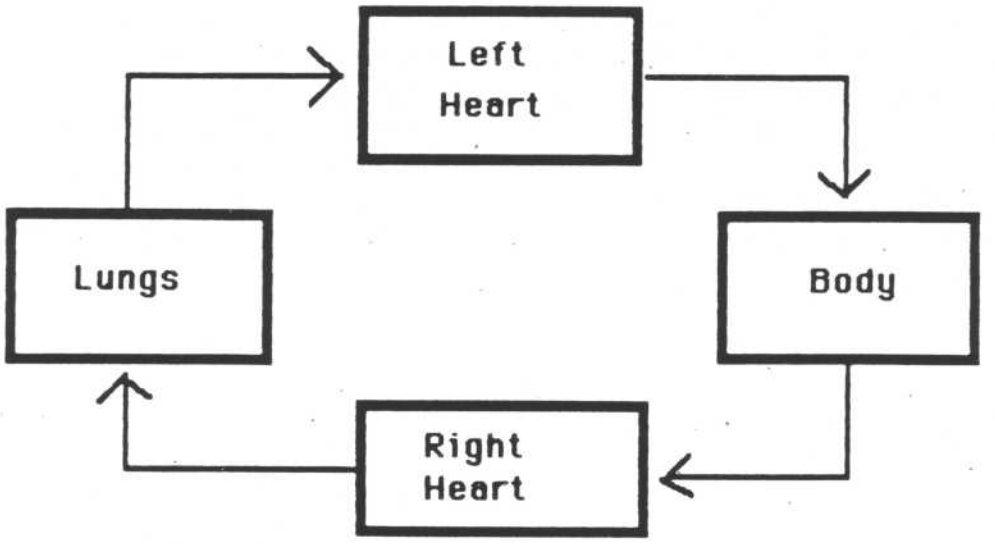

Figure 2: Abstract Circulatory Loop

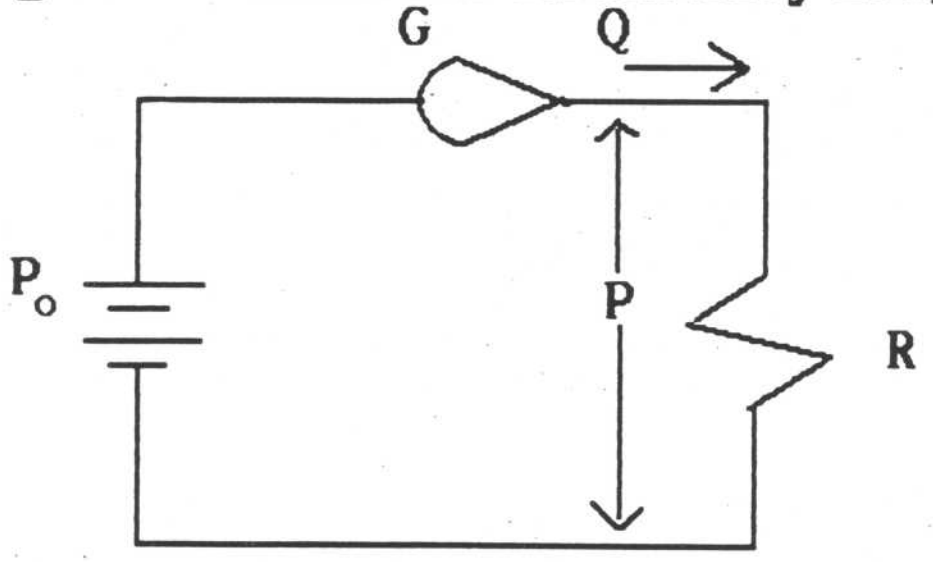

$$Q = G(P_o - P)$$

$$P = RQ$$

Variables

Q: Cardiac Output

P: Arterial Pressure

Properties

R: Total Circulatory Resistance

G: Heart Contractility

P_o: Maximal Heart Pressure

and pulmonary loops varies inversely with venous compliance. A highly compliant vein stretches or sags to accomodate more blood without drastically raising its pressure, thus functioning like an electrical capacitor. Because active-blood volume is the single most important factor in circulatory behavior, and only the veins have dynamic compliances (regulated by the body), venous compliance is a crucial property whose changes incur inverse changes to all pressure and flow variables in the circulatory system.

Using the standard mapping of pressure to voltage and flow to current, Figure 2 (Campbell, 1983) abstracts the circulatory loop (either systemic or pulmonary) into a simple electrical model. A flow source of maximal internal pressure, P_0, and contractile strength, G, outputs flow, Q, against pressure, P, induced by resistance, R. Intuitively, Q varies directly with the pressure differential, (P_0 - P), and with pump (heart) strength. As a simple diagnostic example, if P increases while Q decreases, then P = QR indicates that R must have increased. This corresponds to a doctor's diagnosis of a clogged artery, a common source of super-normal resistance, to account for depressed cardiac output and elevated arterial pressure.

3 Qualitative Sensitivity Analysis

Due to the presence of feedback, via both the cyclic flow of blood and the bi-directionality of component interactions, cardiovascular variables are sensitive to changes in many properties. Hence, a great many implicit constraints relating single properties to single variables underlie the causal model. By uncovering these associations, many of which are non-local, the diagnostician can circumvent causal reasoning and take advantage of the highly-constrained model to identify faults after minimal testing. Campbell (1983) introduces quantitative sensitivities to express the dependence of variables on properties. Calculated as the ratio of partial differentials, the sensitivity of P to R, for instance, is:

$$(\partial P/P)/(\partial R/R) = 1/(1 + G * R) \tag{1}$$

This represents the system-wide sensitivity of P to changes in R under the single-fault assumption that no other properties have changed. Sensitivities provide useful diagnostic pointers from changing variables (i.e. symptoms) to their most strongly-coupled properties, where a "strong" sensitivity has an absolute value close to (but never above) 1. However, during most diagnostic reasoning, the salient aspect of any sensitivity is its sign. Does the property affect the variable directly (positive sensitivity value), inversely (negative sensitivity value) or not at all (sensitivity close to 0)? In short, a good deal of diagnostic reasoning exploits only qualitative information: "If the arterial pressure is up, then the vascular compliance might be down, or the arterial resistance might be up." Quantitative sensitivities exceed informational needs while charging a large computational cost.

We propose qualitative sensitivity analysis (QUALSA) to extract only the necessary diagnostic information from a set of constraints. QUALSA begins by converting all system constraints to "mixed" confluences (De Kleer and Brown, 1985). Next, a set of parameter assumptions is created to define a qualitative state of the system. These assumptions then enable a one-to-one mapping from mixed confluences to "pure" confluences. The latter characterize the dynamic behavior of the qualitative state. Now, to test the system-wide qualitative sensitivities of all variables to a selected property ϕ, $\partial\phi$ is set to either + or -, and all other property' derivatives are set to 0. By restricting the values of all confluence terms to the (-,0,+) quantity space (Forbus, 1985), QUALSA encounters the ambiguities of qualitative arithmetic (Simmons, 1986). Using a constraint-satisfaction technique capable of dealing with these ambiguities (Thyagaragan, 1987), all valid interpretations of the qualitative state's confluence set are found subject to the previous assignment of property-derivative values. The collection of qualitative variable values from each interpretation serves as a fault-table index for $\partial\phi$ = x, where x is the original setting of $\partial\phi$. For each interpretation, a comparison of ∂X to $\partial\phi$ for any variable X yields a qualitative sensitivity defined as:

$$QLS(X,\phi) = \begin{cases} + & \text{if } \partial X = \partial \phi \\ - & \text{if } \partial X \neq \partial \phi \text{ and } \partial X \neq 0 \\ 0 & \text{if } \partial X = 0 \end{cases} \quad (2)$$

4 Applying the QUALSA Algorithm

When applied to the abstract cardiovascular model of Figure 2, QUALSA proceeds as follows:

I. *Initial equations with properties = (G, P_0, R) and variables = (P, Q)*

$$Q = G * (P_0 - P) \quad (3)$$

$$P = R * Q \quad (4)$$

II. *Differentiate and transform to mixed confluences. [x] represents the sign of x, whether +, - or 0.*

$$\partial Q = \partial G * [P_0 - P] + [G] * (\partial P_0 - \partial P) \quad (5)$$

$$\partial P = \partial R * [Q] + [R] * \partial Q \quad (6)$$

III. *Make parameter assumptions:*

$$P_0 > P \ , G > 0 \ , Q > 0 \ , R > 0$$

IV. *Apply parameter assumptions to mixed confluences to derive pure confluences:*

$$\partial Q = \partial G + \partial P_0 - \partial P \quad (7)$$

$$\partial P = \partial R + \partial Q \quad (8)$$

V. *Modify a property (plant a fault):*

$$\partial G \leftarrow +$$

VI. *Apply Single-Fault Assumption:*

$$\partial P_0 \leftarrow 0 \ , \partial R \leftarrow 0$$

VII. *Call constraint satisfier with confluences 7 and 8 and instantiated properties. Receive a unique valid interpretation:*

$$(\partial Q+, \partial P+)$$

VIII. *Calculate Qualitative Sensitivities of both variables to G using Definition 2:*

$$(\partial P = \partial G) \Rightarrow (QLS(P,G) \leftarrow +)$$

$$(\partial Q = \partial G) \Rightarrow (QLS(Q,G) \leftarrow +)$$

IX. *Repeated calls to the constraint satisfier under the single-fault assumption with each property faulted high and low yield a complete fault table:*

	$\partial P+$	$\partial P0$	$\partial P-$
$\partial Q+$	$\partial G+$ or ∂P_0+	nil	$\partial R-$
$\partial Q0$	nil	nil	nil
$\partial Q-$	$\partial R+$	nil	$\partial G-$ or ∂P_0-

Table 1: Faults indexed by symptoms

X. *Calculate all qualitative sensitivities , which in this case remain unambiguous over the 6 interpretations returned by the 6 calls to the constraint satisfier:*

	G	Po	R
P	+	+	+
Q	+	+	-

Table 2: Qualitative Sensitivities

5 Ambiguities

The ambiguity of qualitative arithmetic often spurns an abundance of interpretations, any two of which will have conflicting sensitivities QLS(X, ϕ ,I1) and QLS(X, ϕ,I2) for at least one variable X and property ϕ. Hence, qualitative sensitivities are sometimes indeterminate. In addition, selected properties ϕ_1 and ϕ_2 can yield one or more of the same interpretations when individually instantiated to + or - and run through QUALSA. This creates a one-to-many mapping of indices to faults in the fault table. Multiple interpretations for a single property setting (fault) contribute to a many-to-one mapping, but this creates no additional ambiguity for the backward causal reasoning indigenous to diagnosis.

6 Aggregation

The sensitivities of Table 2 indicate that G and P_0 have identical qualitative effects upon variables P and Q. In fact, a more detailed cardiovascular model reveals further similarities in their induced sensitivities. These similarities, along with their common location, the left heart, make P_0 and G excellent candidates for a simplifying aggregation. Let H represent a general heart strength and define it as:

$$H = G * P_0 \tag{9}$$

Under the assumptions:

$$G > 0, P_0 > 0$$

steps II-IV of QUALSA produce the pure confluence:

$$\partial H = \partial G + \partial P_0 \tag{10}$$

Substituting Equation 10 into Equation 7, a legal substitution in qualitative arithmetic (De Kleer and Brown, 1986) since the coefficient of $G * P_0$ is the same in both Equation 9 and 3 , yields:

$$\partial Q = \partial H - \partial P \tag{11}$$

Four applications (two for each property, H and R) of QUALSA steps V-VIII to confluence Equations 11 and 8 generate simplified fault and sensitivity tables:

	$\partial P+$	$\partial P0$	$\partial P-$
$\partial Q+$	$\partial H+$	nil	$\partial R-$
$\partial Q0$	nil	nil	nil
$\partial Q-$	$\partial R+$	nil	$\partial H-$

Table 3: Aggregated Fault Table

	H	R
P	+	+
Q	+	-

Table 4: Aggregated Sensitivity Table

Now, diagnosis can proceed in the abstracted fault space containing only H and R. Only if the fault is localized to H will granularity shift to the level of P_0 and G, where variables such as the left heart's diastolic("filling") and systolic("emptying") blood volumes will discriminate between the two primitive faults.

7 Multiple Faults

Under the strong assumption that no confluence contains more than a single property derivative, a "comprehensive" multiple-fault table (i.e. one that covers from 0 to n faults, where n is the number of model' properties) can be efficiently generated. Normally, this would require 3^n calls to the constraint satisfier to test the affects of all combinations of property-derivative settings (+,-, and 0). But under the single-property assumption, only 3n such calls are required, where each call involves only the confluences containing a specified property. The interpretations returned from each such call are then intersected with the interpretations from other calls to create indices into the comprehensive fault table.

For example, Equations 11 and 8 have only single property derivatives. Setting:

$$\partial H \leftarrow +$$

and passing Equation 11 to the constraint satisfier creates interpretation set:

$$ISET1 = ((\partial P+, \partial Q-), (\partial P+, \partial Q0), (\partial P+, \partial Q+), (\partial P0, \partial Q+), (\partial P-, \partial Q+)) \quad (12)$$

Next, let:

$$\partial R \leftarrow -$$

and pass Equation 4 to the constraint satisfier. This returns:

$$ISET2 = ((\partial P-, \partial Q-), (\partial P-, \partial Q0), (\partial P-, \partial Q+), (\partial P0, \partial Q+), (\partial P+, \partial Q+)) \quad (13)$$

After intersecting ISET1 and ISET2 to yield:

$$ISET3 = ((\partial P-, \partial Q+), (\partial P0, \partial Q+), (\partial P+, \partial Q+)) \quad (14)$$

use each of the three ISET3 interpretations as an index for the double fault ($\partial H+$ and $\partial R-$).

8 Conclusion

The nature of complex systems precludes the use of only local behavioral knowledge to diagnose faults efficiently. QUALSA exploits local behavioral constraints to uncover implicit interactions, both local and global, between properties and variables. Its qualitative basis avoids the algebraic complexities of quantitative methods at the cost of increased ambiguity in both the forward causal reasoning indigenous to simulation and in backward diagnostic reasoning. By exploiting empirical ordinal relationships, such as that fact that arterial pressure normally greatly exceeds venous pressure, and organizing them in a quantity lattice (Simmons, 1986) or similar structure, we expect to reduce this nondeterminism considerably.

Campbell (1983) has detailed the drastic sensitivity alterations incurred by minor modifications to the component topology, while de Kleer and Brown (1983) have illustrated the importance of locality (and more generally, no function in structure) for robust modelling. Thus, sensitivities can discriminate among the behaviors of structurally similar models and thereby capture the behavioral ramifications of minor structural adjustments; and their derivation from local constraints enhances robustness. In short, sensitivities are well suited for models of evolving systems.

In addition to supporting structural changes, sensitivities can suggest structural abstractions to strengthen diagnostic capabilities. These aggregations embody a more organized understanding of the modeled system – an understanding recognized and implemented in diagnostic systems such as ABEL(Patil et. al, 1982) and INTERNIST-II (Pople, 1982) but supplied externally. By deriving structure from within, the integration of QUALSA and aggregation exhibits a theory of self-contained diagnostic learning that unites experiential and first-principle techniques for their mutual enrichment.

9 References

- Campbell, Kenneth, The Physical Basis of the Problem Environment in Cardiovascular Diagnosis, January, 1983, unpublished.
- De Kleer, J. and Brown, J.S., Assumptions and Ambiguities in Mechanistic Mental Models , in: D. Gentner and A.S. Stevens (Eds.) *Mental Models* (Erlbaum, Hillsdale, NJ, 1983) 155 - 190.
- De Kleer, J. and Brown, J.S., Theories of Causal Ordering, *Artificial Intelligence* 29(1), 1986.
- Forbus, K.D., Qualitative Reasoning About Physical Processes, *Artificial Intelligence* 24, 1984.
- Peterson, Nils S. and Campbell, Kenneth B., Teaching Cardiovascular Integrations with Computer Laboratories ,*The Physiologist* , Vol. 28, No. 3, 1985.
- Patil, Ramesh, Peter Szolovits and William Schwartz, Modeling Knowledge of the Patient in Acid-Base and Electrolyte Disorders, in *Artificial Intelligence in Medicine*, ed. Peter Szolovits, Westview Press, 1982.
- Pople, Harry E. Jr., Heuristic Methods for Imposing Structure on Ill-structured Problems: The Structuring of Medical Diagnosis, in *Artificial Intelligence in Medicine, ed. Peter Szolovits*, Westview Press, 1982.
- Reiter, Raymond, A Theory of Diagnosis from First Principles, Univ. of Toronto, Tech Rep. No.187/86, to appear in *Artificial Intelligence*
- Simmons, Reid, "Commonsense" Arithmetic Reasoning, *Proceedings of the Fifth National Conference on Artificial Intelligence* , 1986.
- Thyagarajan, P.,forthcoming master's thesis, The Univ. of Oregon, 1987, unpublished.
- Weld, D.S. Use of Aggregation in Causal Simulation, *Artificial Intelligence* , 30(1), 1986.