1

REASONING ABOUT THERAPY FROM A PHYSIOLOGICAL MODEL

W. J. Long, S. Naimi, M. G. Criscitiello, S. Kurzrok

M.I.T. Laboratory for Computer Science, Cambridge and
Tufts-New England Medical Center, Boston, Massachusetts

Biomedical Measurement Informatics and Control, **2:** 40-45, 1988

Abstract

In many medical domains a causal physiological model provides a knowledge base of relationships useful for reasoning about diagnosis, management, prognosis, and basic understanding of the processes. This paper reports on our efforts to develop algorithms for reasoning about the potential effects of therapy in the context of a system designed to assist the physician in diagnosis and management of patients with heart failure. The primary result has been the development of an algorithm based on signal flow analysis for predicting effects and the implementation of this algorithm as a mechanism for handling multiple effects, changes over time, non-linear relationships, and in providing explanations. When applied to a model of the cardiovascular system we are developing this methodology predicts drug effects consistent with the medical literature.

1 INTRODUCTION

The purpose of this research is the development of a program to aid the physician with the diagnosis and management of patients with cardiovascu-

[0]This research was supported by National Institutes of Health Grant No. R01 HL33041 from the National Heart, Lung, and Blood Institute and No. R24 RR01320 from the Division of Research Resources.

lar disease characterized by heart failure (cardiac performance inadequate to meet the body's demand). The domain is challenging because:

1. Physiological reasoning is important to determine the actual cause among many potentially responsible for producing the compensatory and non-compensatory responses that characterize heart failure.

2. Many causes are not correctable, making diagnosis and management of aggravating factors and the interrelationship of management and patient state important.

3. Different aspects of the system change at different rates. While many parameters, such as heart rate and blood pressure change rapidly in response to changes, other factors such as blood volume and heart size change more slowly.

Our ultimate goal is to provide a tool that the practitioner can use as a *reasoning network* for thinking about the state of the patient, the potential benefit of more information, the implications of hypotheses, the possibilities that have been ruled in or out, and the likely effects of therapy.

The program is organized around a qualitative causal physiological model of the cardiovascular system. This model serves as the knowledge base specifying the relationships among parameters as well as the central repository for patient specific information. The knowledge base is used to interpret patient input in terms of model parameter states, to reason about possible causes and complications, to look for therapies and to assess their potential. The physiological model initially represents the knowledge about relationships between physiological, therapeutic, and primary causal parameters. As reasoning takes place the model is constrained to represent what is known about the patient.

Reasoning starts when the input module interprets the clinical data as evidence constraining the model parameters. Relationships between parameter values, such as possible and necessary cause and effect relationships, are supported by a truth maintenance system[1] that propagates the logical implications of the parameter values. The diagnostic module attempts to trace undesirable effects (pulmonary edema, fatigue, poor renal function, angina, etc.) to their ultimate causes and possible aggravating factors. The therapeutic module uses these causal chains to identify therapies that may break the chains. Since therapies may have multiple effects or a given effect may have multiple implications in the overall system, the therapies may have effects that will worsen the patient state. The program must analyze the potential effects to assist the physician in anticipating both the expected and unexpected outcomes. Additionally, it is useful to know the principal factors contributing to potential changes justifying the reasoning and alerting the physician to the critical assumptions about the patient.

This paper concentrates on the methodology, implementation, and effectiveness of an algorithm for predicting the effect of therapy. A more thorough overview of the rest of the program appears elsewhere[2].

2 PREDICTING CHANGES

Reasoning about the effects of a therapeutic intervention is analogous to reasoning about the effects of a perturbation on a network of constraints and influences. The intuitively appealing strategy is to follow the perturbation through the network drawing qualitative or quantitative conclusions about the additional influence of individual relationships, reasoning about combinations of influences where divergent paths in the network rejoin, and following the changes until the ripples die out. Such reasoning needs to consider the nature of the relationships in the networks, whether the parameters represent levels or rates, whether the combinations are additive or multiplicative, the effects of feedback, and so forth. The ultimate realization of such a strategy is the simulation of a network, such as Guyton's cardiovascular simulation[3]. The major difficulty with such an approach is the need to know the current value of each parameter in the patient to start the simulation. Given that detailed information, the model computes a behavior consistent with current physiological knowledge. Since many of the parameters can only be measured invasively if at all, there is no way to determine all of the necessary data in the individual. A second strategy is to make the simulation model qualitative using only the direction and nature of relationships. Kuipers has shown that it is possible to do useful reasoning about the qualitative behavior of simple systems with such a methodology[4]. The problem in the cardiovascular domain is the explosion of possible model states when adding opposing influences of unknown magnitudes. As a result, the predictive value of the model is sacrificed. A general problem with all simulation approaches is the difficulty of explaining the simulation results. The importance of different pathways is lost in the computation, leaving one with only the predicted parameter changes. The user really needs to know the sensitivity of the results to different factors to apply his own more detailed knowledge of the patient and assess the importance of any assumptions he may have made.

The approach we have chosen is different. By assuming the system will reach a stable state after a perturbation, the question becomes how that stable state is changed from the state before therapy. Signal flow analysis takes that same approach in the domain of circuit analysis to predict the circuit gain. With some simplifying assumptions and modifications to the signal flow analysis machinery, we are able to apply the approach to reasoning about a physiological model. The assumptions are as follows:

1. The system goes from steady state to steady state. The justification is that the cardiovascular system is highly stable[5]. However, there are several time periods over which different parts of the system reach stability. We assume that for the time period of concern, parts with shorter time constants have reached a stable state and parts with longer time constants have no effect.

2. The system can be modeled as piece-wise linear. The major obstacles to linearity are non-linear relationships between parameters and multiplicative influences on a given parameter. When two parameters are related non-linearly (*e.g.,* the Frank-Starling curve between end diastolic pressure and stroke volume), the relationship can be divided into nearly linear regions. When a parameter is the product of two others, the relationship of the changes is:

$$\triangle(A \times B) = \triangle A \times B + \triangle B \times A + \triangle A \times \triangle B$$

   Thus, if the product of changes term is dropped, the relationship again becomes piece-wise linear.

Given these simplifications, the network of physiological parameters is linear and the techniques of signal flow analysis become applicable. In particular, we could use Mason's General Gain formula[6] to determine the change in any particular parameter, given a change in some other parameter. However, the usual form of this formula is unintuitive and inefficient when all parameter changes are needed.

Hence, we have derived a different formulation that computes the gain incrementally from parameter to parameter, correcting for feedback each time a new feedback path is encountered. (The derivation and implementation of this formula is discussed in[7].) Essentially, the computation involves computing for each path from the changed parameter to other parameters the gain for each link. The gain is the inherent gain of the link adjusted for any new feedback loops encountered by the path at that point and the change in a parameter is the sum of the path gains going through the parameter.

3 CARDIOVASCULAR MODEL

We are developing a cardiovascular model to meet the requirements of the Heart Failure Project and the assumptions of the algorithm discussed in the previous section. Several cardiovascular simulation models are reported in the literature, but none quite fits our needs. The parameters should be ones of concern to the cardiologist and, as much as possible, measurable in the patient. Also, the model must include factors of importance in the kinds of heart disease within the domain without including extraneous factors. Finally, where there are choices in the topology of the model because of multiple constraints, the form most useful for explanation should be chosen. Models such as HUMAN[8] include parameters such as the mean systemic pressure, which is useful for defining the mathematical relationships for venous return, but is not easily measurable in vivo. That model also contains relationships for electrolyte behavior and divisions of the circulation into regional flows that are unimportant for our domain. Furthermore, it does not include appropriate parameters to capture conditions such as valve disease that leads to heart failure.

For these reasons we have chosen to develop our own cardiovascular model, more in tune with the reasoning of the clinical cardiologist. That is, we will use the relationships commonly used by cardiologists wherever possible, rather than those used by physiologists. The model as presently constructed (explaining changes described later) is in figure 1. The model includes parameters to account separately for the function of the right and left sides of the heart, the determinants of myocardial ischemia, the sympathetic reflexes, and the renal determination of blood volume. The parameters for each side of the heart include atrial pressure, end diastolic pressure and volume, compliance, systolic function, systolic pressure, and stroke volume. These parameters are sufficient to represent the primary abnormalities of function, valves, compliance, and dilatation. The sympathetic state is divided into alpha and beta states to represent the different effects that specific sympathetic agents have.

The relationships on the links between parameters are formulas determining the link strength from the parameter states that influence it. If the relationship is linear, this is just the strength of the link. The link strengths in the current model can be zero (decoupled) or positive or negative with value either 0.5, 1.0, or 1.5. (The algorithm could support any values.) The parameter values themselves are scaled such that these link strengths represent weak, moderate, and strong relationships. These values have proven sufficient to represent the experiential knowledge of the relations. As demonstrated in the next section, they are also sufficient to account for the behavior of the system reported in the literature.

The model is still in active development, so we do not claim this is the final form. Our refinement strategy is to add parameters only as necessary to account for effects of pertinent disease states or therapies or to clarify the explanations.

4 APPLYING CHANGE ANALYSIS TO THE MODEL

Use of the signal flow analysis algorithm places several requirements on the implementation of the model and program. The algorithm must run fast enough to reasonably consider different therapies in one session. The implementation must allow for effective explanations of the results. The implementation must support reasoning about parameters that change across non-linear regions, about changes over longer periods of time, and with uncertain parameter states.

The signal flow analysis algorithm is implemented in stages. All feedback loops are computed upon loading the model. That allows the path generation to determine the loops encountered by the path at each new parameter. The gain along a path is the product of the gains across the links and the total gain at any parameter is then the sum of path gains to the parameter. Multiple changes to the system, such as drug combinations or drugs that act on more than one physiological parameter are handled by summing the changes. The loops and parameters are both represented as bit vectors to increase the efficiency of the many membership comparison operations. Together, these techniques increase the efficiency such that the computation of the changes to all parameters presently from a single parameter requires a few seconds on a Symbolics 3640 — an acceptable speed.

The implementation enables an effective explanation technique. Since parameter changes are determined by summing the changes along the various pathways, the contributions of the pathways can be compared. To explain a parameter change, the program examines the contributing pathways and highlights the pathway making the largest contribution and those making some threshold as much. The algorithm is actually more complicated to handle compensating effects. For example, heart rate affects cardiac output both directly and by changing the stroke volume. Normally the two effects virtually cancel each other. Hence, when both heart rate and cardiac output change, the highlighting algorithm may highlight these two paths unless some addition is made to eliminate such sets of paths. This is accomplished by collecting the paths by sets that touch the same feedback loops and only highlighting paths from those sets that together contribute to the change. The result is the very graphic way of displaying the primary mechanisms that determine the decreased likelihood of myocardial ischemia produced by a beta-blocker in figure 1.

Since the changes caused by therapies are not small — after all, the whole point is to have a substantial corrective effect — the piece-wise linear approximations must be handled. The program determines which link gains will change first from the changes and the current parameter values. These are then changed to the next region and changes are recomputed. In this way the total response can be computed. This method is not guaranteed to make the appropriate changes since the transient behavior may be different from the stable state, but in a highly damped system such as the cardiovascular system the assumption is reasonable.

Another problem is that different parts of the cardiovascular system take different amounts of time to stabilize. That is, the short-term effects of a therapy may be different from the long-terms effects. The program assumes that pathways with long time constants have no effect on short-term solutions and changes are determined separately for the different time periods. Thus, the algorithm for determining the changes from large dosages of drug over a long period of time is: Starting with the shortest time constants in the system, determine the changes for a small dosage. Change the parameter values affecting non-linear gains and recompute as necessary to determine the immediate changes for the appropriate size dosage. Move to the next time period with the projected parameter values and compute the changes including links that have effects within that time period. Continue until the desired predictions are determined.

Figure 1: Cardiovascular model showing effects of a beta-blocker

5 APPLICATION OF THE MODEL TO THERAPY

To validate the approach, we compared predictions of the model in the normal state with information from medical literature on the effects of the major classes of drugs used for the treatment of patients with heart failure and coronary artery disease. The drugs are represented by adding each as a model parameter affecting those parameters *directly* affected by the drug. Some drugs have a single effect and therefore link to a single parameter. Propranolol, as a typical beta adrenergic blocker decreases the beta sympathetic state. Other drugs affect multiple parameters. For example, nitroglycerin primarily causes venodilation (increased venous compliance), but also causes some amount of vasodilation (decreased systemic resistance) as well as some reports of preferential coronary artery dilation. Thus, the direct influence on systemic resistance is sufficiently negative to account for the reported overall unchanged systemic resistance. The drugs selected and their direct effects are in figure 2.

Figure 3 shows the changes reported in the literature versus the results computed from the model. There is variation in the literature on the effects, so we used a single recent review article as the source wherever possible[9]. This article only reported direction and not relative amount, however we have included relative amounts in the model predictions. These model predictions are normalized so that the largest change for each drug is represented as three arrows and smaller changes as lesser numbers of arrows. Hence, relative strengths of drugs are not indicated by the arrows. The

drug	drug type	direct influence
hydralazine	vasodilator	↓systemic resistance
propranolol	beta blocker	↓beta state
furosemide	diuretic	↑venous compliance
dobutamine	inotropic agent	↑inotropic state and .3 ↑heart rate
nitroglycerin	venodilator	↑venous compliance and .1 ↓systemic resistance
nifedipine	calcium channel blocker	↓systemic resistance and .3 ↓inotropic state
verapamil	calcium channel blocker	↓systemic resistance, .8 ↓inotropic state, and .8 ↓heart rate

Figure 2: Direct influences of representative drugs

effects of nitroglycerin and furosemide are not included in this article and come from Goodman and Gilman[10].

The model predictions are mostly consistent with the literature. The three differences were the change of hydralazine and nifedipine on LVEDP and that of verapamil on CO. The predictions for the verapamil are consistent with other reports in the literature and the LVEDP changes are small and would probably not be there if the patients tested were actually in heart failure. Overall the model effectively predicts the outcomes of the therapies.

The model is also capable of predicting effects of therapy in disease states. In figure 4, aortic stenosis (AS) is represented as a disease state with a strong relationship between LVEDP and stroke volume (more dependent on filling pressure) and a weak relationship between LV systolic function and LV emptying (little response to increased inotropy). The predictions for three drugs in the normal are compared to the effects in AS. There is no relevant literature to compare these predictions to, but clinically they make sense. The model predicts that in AS the heart is less responsive in increasing cardiac output with decreases in blood pressure or increases in inotropy, and more vulnerable to fall in cardiac output and blood pressure with reduced LVEDP.

drug		HR	AP	SVR	CO	EDP
hydralazine	data	↑	↓	↓	↑	→
	model	↑↑	↓↓	↓↓↓	↑↑	↓
propranolol	data	↓	↓	↑	↓	↑
	model	↓↓↓	↓	↑	↓↓	↑
furosemide	data					↓
	model	↑	↓	↑	↓	↓↓↓
dobutamine	data	↑,→	↑,→	↓	↑	↓
	model	→	↑↑	↓↓	↑↑↑	↓↓
nitroglycerin	data	↑,→	↓	→	↓	↓
	model	↑	↓	→	↓	↓↓↓
nifedipine	data	↑	↓	↓	↑	→
	model	↑↑	↓↓	↓↓↓	↑	↓
verapamil	data	↓,→	↓	↓	↑	↑,→
	model	↓	↓↓↓	↓↓↓	→	→

HR = heart rate, AP = arterial pressure, SVR = systemic vascular resistance, CO = cardiac output, EDP = left ventricular end diastolic pressure

Figure 3: Comparison of predictions to literature

drug		HR	AP	SVR	CO	EDP
hydralazine	normal	↑↑	↓↓	↓↓↓	↑↑	↓
	AS	↑↑↑	↓↓↓	↓↓↓	↑	↓
dobutamine	normal	→	↑↑	↓↓	↑↑↑	↓↓
	AS	↑	↑	↓	↑	↓
nitroglycerin	normal	↑	↓	→	↓	↓↓↓
	AS	↑↑	↓↓	↑	↓↓↓	↓↓↓

Figure 4: Predictions for normal and aortic stenosis

References

[1] McAllester, David A., "Reasoning Utility Package User's Manual, Version One," MIT/AIM–667, (Massachusetts Institute of Technology, Cambridge, 1982).

[2] Long, W. J.; Naimi, S.; Criscitiello, M. G.; Pauker, S. G.; and Szolovits, P., "An Aid to Physiological Reasoning in the Management of Cardiovascular Disease," *1984 Computers in Cardiology Conference*, (1984) 3-6.

[3] Guyton, A. C.; Coleman, T. G.; and Granger, H. J., "Circulation: Overall Regulation," *Annual Review of Physiology*, (1972) **34:** 13-46.

[4] Kuipers, Benjamin and Kassirer, Jerome P., "Causal Reasoning in Medicine: Analysis of a Protocol," *Cognitive Science*, (1984) **8:** 363-385.

[5] Guyton, Arthur C.; Coleman, Thomas G.; Manning Jr., R. Davis; and Hall, John E., "Some Problems and Solutions for Modeling Overall Cardiovascular Regulation," *Mathematical Biosciences*, (1984) **72:** 141-155.

[6] Mason, Samuel J. "Feedback Theory — Further Properties of Signal Graphs," *Proceedings of the IRE*, (1956) **44:** 920-926.

[7] Long,W. J.; Naimi, S.; Criscitiello, M. G.; and Kurzrok, S., "Predicting the Effects of a Therapy in a Physiological Network," to be published.

[8] Coleman, Thomas G. and Randall, James E. "HUMAN: A Comprehensive Physiological Model," *The Physiologist*, (1983) **26:** 15-21.

[9] Nestico, Pasquale F.; Hakki, A-Hamid; and Iskandrian, Abdulmassih S., "Effects of Cardiac medications on ventricular performance: Emphasis on evaluation with radionuclide angiography," *American Heart Journal*, (1985) **109:** 1070-1084.

[10] Gilman, Alfred G.; Goodman, Louis S.; Rall, Theodore W. and Murad, Ferid, *The Pharmacological Basic of Therapeutics*, (Macmillan, New York, 1985).