


MIT/LCS/TM-240

A PROGRAM FOR THERAPY
OF
ACID-BASE AND ELECTROLYTE DISORDERS



Hank Bromley

June 1983

A Program for Therapy of Acid-Base and Electrolyte Disorders

by
Hank Bromley

Submitted to the Department of Electrical Engineering and Computer Science
on May 20, 1983
in partial fulfillment of the requirements for the degree of Bachelor of Science

Abstract

This thesis describes work done on the therapy component of an on-going project for the diagnosis and management of acid-base and electrolyte disorders. Therapeutic interventions can be classified as *symptomatic* or *etiologic*, and as *acute* or *chronic*. We have focused on the problem of acute symptomatic therapy. Based on observation of clinical practice, we have developed a formalization of the domain-independent aspects of the task of acute symptomatic therapy, then applied the formalization to the particular field of acid-base and electrolyte disorders. A rule-based program named ABET (the Acid-Base and Electrolyte Therapy Advisor) has been designed and written to test this formalization.

The thesis presents the methods used by ABET, the program's implementation, a sample session, and a discussion of limitations and possible improvements.

Thesis supervisor: Ramesh Patil
Title: Research Associate, Laboratory for Computer Science

This research was supported (in part) by NIH Grant No. 1 P01 LM 03374-03 from the National Library of Medicine.

Keywords: AI in Medicine, Expert Systems, Therapy

Acknowledgments

I am gratefully indebted to the following:

Peter Szolovits, for first involving me with this research, and continuing to provide timely advice on matters both technical and otherwise.

William B. Schwartz, for encouraging a neophyte's efforts to comprehend the medicine of acid-base and electrolyte balance, and, together with Prof. Szolovits, supplying the conceptual framework for my project.

Bill Long, for his extremely valuable observations and comments on my work and writings.

Tom Russ, for all kinds of encouragement and assistance, especially his tutelage in the use of the R text formatter.

Glenn Burke, for somehow keeping the ML computer from disintegrating entirely, and meanwhile showing me some simple essentials on how to get it to do what I want.

Elisha Sacks, for reading and making suggestions on a draft of this thesis.

everyone on the 3rd floor, for providing a friendly and stimulating place to work.

and

Larry Widman, for his extraordinarily active collaboration: enthusiastically assuming the frustrating chore of teaching a simple program some complicated medicine; sharing ideas on both overall strategies and specific programming techniques; and demonstrating a genuine interest in the well-being of the programmer as well as the program.

finally, Ramesh Patil, my supervisor, for guidance on all aspects of thesis research. I was indeed fortunate to have a supervisor who is knowledgeable in the theory and practice of Expert Systems, the medicine of acid-base and electrolyte balance, and even the mechanics of producing a polished document, and who - even more fortunately - positively enjoys explaining what he knows. It pleases me to have been his first thesis advisee, and I'm delighted that he stands to have many more.

CONTENTS

1. Introduction	5
1.1 History and Overview of the Acid-Base and Electrolyte Project	5
1.2 The Therapy Component	6
2. The Task Domain	8
2.1 Theory of Symptomatic Therapy	8
2.1.1 The Clinical Practice of Therapy	8
2.1.2 Formalization of the Practice of Therapy	9
2.2 Particular Considerations for Acid-Base and Electrolyte Therapy	11
3. The Algorithm	14
3.1 Initial Recommendations for Individual Abnormalities	16
3.1.1 Determination of Urgency	16
3.1.2 Determination of Initial Treatment Recommendation	17
3.1.3 Determination of Recheck Time	18
3.2 Resolution of Cross-treatment Conflicts	18
3.3 Final Unification of Recommendations	19
3.3.1 Unification of Therapy Modalities	19
3.3.2 Unification of Recheck Times	20
4. The Implementation	20
4.1 Representation of, and Access to, Patient-Specific Knowledge	21
4.2 Actions of Rules	24
4.2.1 Urgency Rules	24
4.2.2 Treatment Rules	25
4.2.3 Recheck Rules	27
4.3 Conflict Resolution (The "Objection Handler")	28
4.4 Final Unification (The "Funnel")	30
5. An Example	31
6. Discussion and Conclusions	40
6.1 Survey of Therapy Programs	40
6.1.1 Mycin	40
6.1.2 CASNET/Glaucoma	41
6.1.3 Digitalis Therapy Advisor	41
6.1.4 VM	42
6.2 Evaluation	43
6.3 Suggestions for Further Work	44
References	47

FIGURES

Fig. 1. The Acid-Base and Electrolyte Consultant System	7
Fig. 2. The space of patient states and the levels of urgency	10
Fig. 3. Henderson-Hasselbalch equation	13
Fig. 4. A schematic for overall system	15
Fig. 5. The patient characteristics	22
Fig. 6. An urgency rule and its application	26
Fig. 7. A treatment rule and its application	27
Fig. 8. A recheck rule	28
Fig. 9. Case Description	32
Fig. 10. Independent Recommendations: pH	33
Fig. 11. Independent Recommendations: sodium	34
Fig. 12. Independent Recommendations: potassium	36
Fig. 13. Conflict Resolution	38
Fig. 14. Unification of Recommendations	39

1. Introduction

Artificial Intelligence research in medical decision-making has been motivated by several goals. First is expanding the availability of quality health care. It is becoming increasingly difficult for physicians to maintain expertise in all the branches of medicine they are likely to run across in their practice. They must rely on the advice of consulting specialists, of whom there are too few. Exacerbating the shortage of expertise is a maldistribution of what resources do exist. Physicians are scarce in rural areas and in inner cities, particularly specialists. Practitioners in these localities are often without adequate expert advice. It is hoped that computer-based expert systems may be developed, capable of serving as consulting specialists where human experts are unavailable.

Second is a possible improvement in the teaching of medicine. Designing programs to simulate the behavior of experts will require a greater understanding of the cognitive processes involved than now exists. Having better models of the reasoning in medical problem solving should allow more effective teaching of this reasoning.

The third motivation for these efforts is the hope of advancing the techniques of Artificial Intelligence itself. The domain of clinical decision-making is a challenging one, both because of its complexity and its strict constraints; the rules governing the behavior of the human body are fixed, and we cannot alter them to render the domain more amenable to existing AI techniques. Should mastery of this field ever be accomplished, it will only be through the development of new methods for using computers to emulate human thought.

1.1 History and Overview of the Acid-Base and Electrolyte Project

With these goals in mind, in 1977 the Clinical Decision-Making Group at MIT's Laboratory for Computer Science, in cooperation with the Tufts New England Medical Center, commenced the implementation of an acid-base and electrolyte consultant system [13]. This field was chosen because it is understood well enough for there to be detailed and sophisticated models of disease and treatment, but not so well understood that the models are complete or easily converted to computational terms. It is a large enough field to provide an ample testing ground for new ideas, yet small enough to allow the construction of a knowledge base in a reasonable amount of time.

The objective of the complete consultant system is to assist in the proper overall management of the patient. The task of *patient management* includes collecting the relevant information, identifying the disease process(es) responsible for the patient's illness, and prescribing treatment to correct the condition. These components of patient management do not have well-defined chronological boundaries. Additional information may be needed to complete the diagnosis, or because the disease itself has evolved over time. Therapy may be used to provide clues to the diagnosis through the patient's response. And the patient's condition, if sufficiently critical, may require therapeutic intervention even before the diagnosis is resolved. The physician thus faces a large range of alternative courses of action. Correspondingly, the program must also be able to choose among the alternatives, maximizing the patient's best interests. From this perspective, the Acid-Base and Electrolyte Consultant system has been designed with separate components for forming a diagnosis, suggesting a therapy, and deciding which alternative to pursue. Modularizing the different components of a physician's knowledge and expertise enables us to evaluate our understanding about each component and their interactions.

A schematic for the entire patient management system is shown in figure 1. It consists of four major components: (1) the Global Decision-Making component, (2) the Diagnosis component, (3) the Therapy component, and (4) the Patient-Specific Model. The patient-specific model represents the state of the patient at any given time. It is the central data structure with which the other components conduct their reasoning.

An implementation of the diagnosis component and the patient-specific model was completed by Dr. Ramesh Patil, in collaboration with Prof. Peter Szolovits and Dr. William B. Schwartz, in 1981 [8]. This thesis describes an initial implementation of the therapy component, the Acid-Base and Electrolyte Therapy Advisor (ABET). The global decision-making component remains to be explored.

1.2 The Therapy Component

Therapy management can be divided into three steps: (a) initial therapy recommendation, (b) therapeutic evaluation, and (c) adaptive reformulation of the therapeutic regimen. We have so far concentrated on step (a), the initial recommendation, leaving for later investigation the evaluation of previous interventions and the refined reformulation of therapy. We have, however, included features to prepare the way for the addition of the latter two.

Our objectives have been to (1) study and formalize the clinical practice of acute symptomatic therapy, (2) explore the feasibility of using existing AI techniques to apply this formalization to the specific field of acid-base and electrolyte disorders, and (3) identify any further mechanisms that may be needed for adequate performance in this domain. Our intent was not to complete a polished, high-performance expert system, but rather to discover what stands in the way of building such a system.

2. The Task Domain

This chapter outlines the task our program was designed to accomplish. The first section is a discussion of therapy in general, including a formalization of observed clinical practice. The subsequent section introduces the requisite medical knowledge specific to the field of acid-base and electrolyte balance, and mentions some particular characteristics of therapy in this field.

2.1 Theory of Symptomatic Therapy

In contrast to the open-ended and difficult to define task of diagnosis, symptomatic therapy involves selection of a specific treatment from the limited array of possibilities at the physician's disposal. The selection criteria are speed and safety. In general, aggressive treatments are most effective but also risky; less aggressive treatments are safer. The physician's job, then, after identifying which conditions are in need of correction, is to choose the particular form of intervention which best balances rapidity and gentleness of effect. Next we consider how clinicians actually pursue this task and present a formalization of the method.

2.1.1 The Clinical Practice of Therapy

Faced with an acute situation, as in emergency rooms and intensive care units, physicians tend to approach the patient by organ system. The functioning of each organ system or subsystem is evaluated, and if found deficient, each abnormality present is further evaluated for the urgency of the need for intervention. The determination of the urgency for an individual abnormality is based on the body's ability to correct the disorder without assistance, and on the hazards of a continuation of the condition; in short, on the likely consequences of failing to intervene. Next a treatment modality consistent with the perceived urgency is selected, and an estimate is made of how long it will be before a noticeable response can be expected. After all the abnormalities have been considered, the collection of interventions is

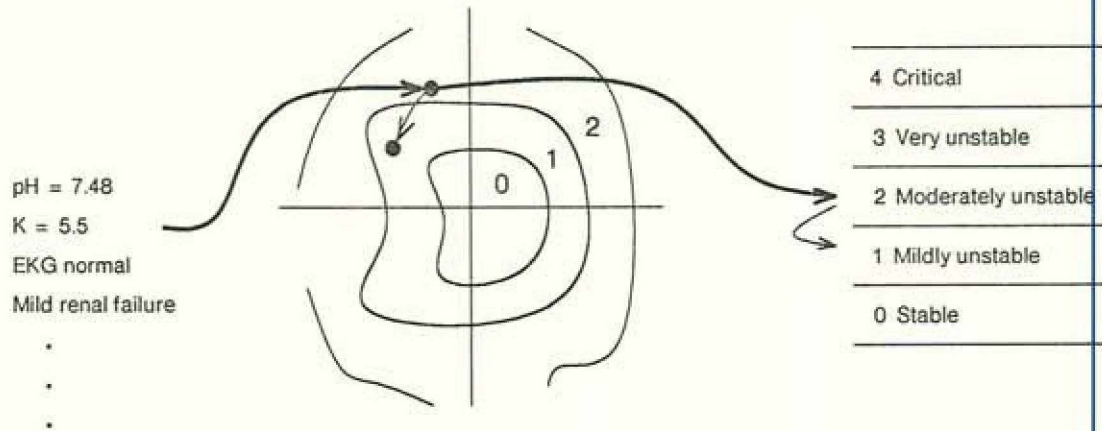
examined as a whole to ensure consistency. When some treatments are found to interfere with others, the lower priority treatments are modified or dropped. A comprehensive treatment plan, correcting the most urgent problems first, and including the expected patient response for comparison against the actual response as a test of the treatment's success, is adopted. Finally detailed plans for routes of administration and follow-up observations are made, taking into account existing forms of intervention and monitoring, such as central venous line access or on-going renal dialysis. As the entire process is conducted in an acute setting, the physician must act swiftly. The methods used for gathering clinical information are therefore geared towards minimizing unnecessary investigation. Classifying the abnormalities by urgency allows the physician to avoid wasting time and cluttering his/her thinking with collecting and weighing unessential information; s/he can concentrate on the facts most pertinent to the stabilization of the patient's condition.

2.1.2 Formalization of the Practice of Therapy

The urgency of the need for treatment clearly plays a central role in the physician's thinking, and thus also dominates our formalization of the therapy process. We metaphorically envision the patient as occupying some point in a space of possible patient states (see figure 2). According to this metaphor, a particular patient is located within the space by mapping from a description of his/her condition onto the axes of the space. The space is divided into regions representing different degrees of stability, ranging from stable in the center (*i.e.*, states which call for no therapeutic intervention) to critical on the periphery. Fully defining the relationship between the description of the patient's condition and the space of patient states is, of course, impossible. We are limited to heuristics assigning regions of the state space to specific patterns of clinical conditions. These heuristics are, in a different language, the same guidelines an examining physician uses to relate observable symptoms to conclusions about a patient's stability.

Once the region containing the patient state is identified, we have an evaluation of the patient's stability, which bears directly on the urgency for therapeutic intervention. The worse the patient's condition is, the more quickly we need to induce an improvement. The goal of any intervention is then to lessen the urgency (bring the stability up) one level (see figure 2). The calculated urgency thus provides a context for all later decisions, specifically including the choice of treatment. Adhering to the goal of moving up one level guarantees a good balance between speed and safety for the therapy ultimately

Fig. 2. The space of patient states and the levels of urgency



recommended, for it ensures that the recommendation will be sufficiently aggressive to induce a significant improvement in the patient's condition, while also being no more risky than it need be to achieve such an effect.

The next step is to, within the context of the perceived urgency, identify the appropriate treatments from a much larger list of possibilities and evaluate each of them. Each has associated with it risk and benefit measures. The particular option with the best combination of risks and benefits, with respect to the goal of reducing the urgency, is chosen.

After this procedure has been completed for each abnormality requiring attention, all the individual recommendations are compared. If any conflicts are discovered, the conflicting recommendations are adjusted to arrive at a resolution, if at all possible.

It should be noted that in the course of formulating these recommendations, the dosages specified need not be precise, because the patient can be reevaluated during the administration of therapy. Further, the initial recommendations often *cannot* be precise because of missing parameters, so it is necessary to adapt therapy to the individual patient by choosing an initial dose and then observing the actual response. If the response is not as expected the discrepancy provides the most useful information

for planning a corrected treatment. For example, in administering bicarbonate therapy, a physician must estimate the space of distribution of bicarbonate in a given patient (bicarbonate distributes throughout a substantial fraction of the total body water). The total body water, however, varies significantly from patient to patient and is often difficult to assess precisely. If the electrolytes are checked after the initial dose has been given, the physician can easily estimate by extrapolation how much more would be required to achieve the therapeutic goals. Any program that recommends therapy must be aware of this very important aspect of clinical practice; it ought to determine not only the dose required but also the waiting time before a reevaluation and possible correction.

These are the principles upon which our work is based. Using them we have been able to construct a program capable of dealing with complex electrolyte disorders, including interactions between physiological systems, in a simple yet general and powerful manner. It is expected that these principles may be applicable to other therapy situations, such as acute management of cardiac arrhythmias, cardiovascular collapse, seizure management and neurological catastrophe, acute respiratory insufficiency and short-term ventilator management, gastrointestinal hemorrhage, multisystem trauma, and immediate poison-control recommendations.

Specifically excluded from the domain of this kind of program, as tasks for which our model of clinical practice is inapplicable, would be diagnosis, intermediate and long-range treatment plans involving sophisticated modeling (e.g., pharmacologic) of patient response, and settings in which therapy is part of the diagnostic procedure. Other projects which have addressed these issues include the Digitalis Therapy Advisor [3], CASNET [15], and the Ventricular Arrhythmia Manager [9].

2.2 Particular Considerations for Acid-Base and Electrolyte Therapy

The following brief discussion of the electrolyte and acid-base disorders should suffice for the reader to understand the description of the program and the examples presented later. This simplified explanation overlooks many of the important subtleties [11], and is not to be taken as a definitive exposition. It should nonetheless be adequate for our purposes.

The water content of the body, approximately 50 to 60 percent by weight, is distributed between the intracellular (within the cells) and extracellular (outside cells) compartments, with the extracellular compartment itself being divided into the interstitial space and the plasma. Water moves freely across


cell boundaries, maintaining osmotic equilibrium between the compartments. The electrolytes, however, due to differences in permeability and active ion pumps, are distributed asymmetrically. The extracellular fluid contains mostly sodium, chloride and bicarbonate, while potassium and organic anions predominate in the intracellular fluid.

The pH of extracellular fluid is normally between 7.35 and 7.45, and is maintained within this range by three mechanisms: (1) the body buffers, (2) pulmonary regulation of CO_2 concentration, and (3) renal (kidney) excretion of acids and alkali. All three act together to handle the normal daily acid load, maintaining a steady-state acid-base equilibrium. As food is oxidized both carbon dioxide (carbonic acid) and nonvolatile acids are added to the extracellular fluid. Immediate buffering minimizes the local pH change, and permits transportation of these acids to the lungs and kidney for excretion. Carbon dioxide is excreted almost entirely by the lungs, and the nonvolatile acids are excreted solely by the kidney. Bicarbonate is regenerated by the kidney as it eliminates the excess acid, replenishing the bicarbonate stores that were originally depleted by the buffering of dietary acid.

The essential carbonic acid - bicarbonate buffer pair is governed by the Henderson-Hasselbalch equation (see figure 3). This equation shows clearly what changes can be expected from both metabolic and respiratory disturbances of acid-base equilibrium. A reduction in bicarbonate concentration will cause the reaction to shift to the right, thus increasing hydrogen ion concentration (*metabolic acidosis*), whereas an elevation in bicarbonate concentration will cause the reaction to shift to the left, thus decreasing hydrogen ion concentration (*metabolic alkalosis*). Similarly, a rise in pCO_2 increases the hydrogen ion concentration by shifting the reaction to the right (*respiratory acidosis*), and a fall has the reverse effect (*respiratory alkalosis*).

As sodium is the ion present in highest concentration in the extracellular fluid, sodium salts are the primary determinant of the osmolality of the extracellular compartment. Changes in sodium concentration consequently have a major influence on the distribution of water between the intracellular and extracellular spaces. Serum sodium is normally stabilized at approximately 140 mEq/L by changes in water balance that occur in response to variations in plasma osmotic pressure. A slight increase in sodium concentration, and in osmotic pressure, leads to the release of antidiuretic hormone and to a retention of water that then restores normal tonicity. Conversely, a slight reduction in serum sodium concentration and osmotic pressure inhibits the release of hormone and permits any excess water to be excreted.

Fig. 3. Henderson-Hasselbalch equation



As the major intracellular cation, potassium also plays a significant part in control of osmotic pressure. Additionally, the potassium concentration of body fluids has an important influence on muscle excitability. Extracellular fluid normally contains 3.8 to 5 mEq of potassium per liter, and intracellular fluid approximately 150 mEq/L; thus only a small fraction of the 2500 to 3000 mEq of potassium in the body is contained in the extracellular space. Therefore relatively small absolute changes in extracellular concentration, by producing large differences in the ratio of intracellular to extracellular potassium, may have important effects on neuromuscular activity.

The potassium distribution is partly determined by the pH. Hydrogen ions move freely across cell membranes, and are thus evenly distributed throughout the body fluids. Any increase in the bloodstream's H^+ concentration (decrease in serum-pH) is accompanied by an identical change in the intracellular fluid. As H^+ enters the cells, to prevent a net positive charge within the cells, an equivalent amount of K^+ must be dumped out, increasing the serum-k. Similarly, a drop in the serum

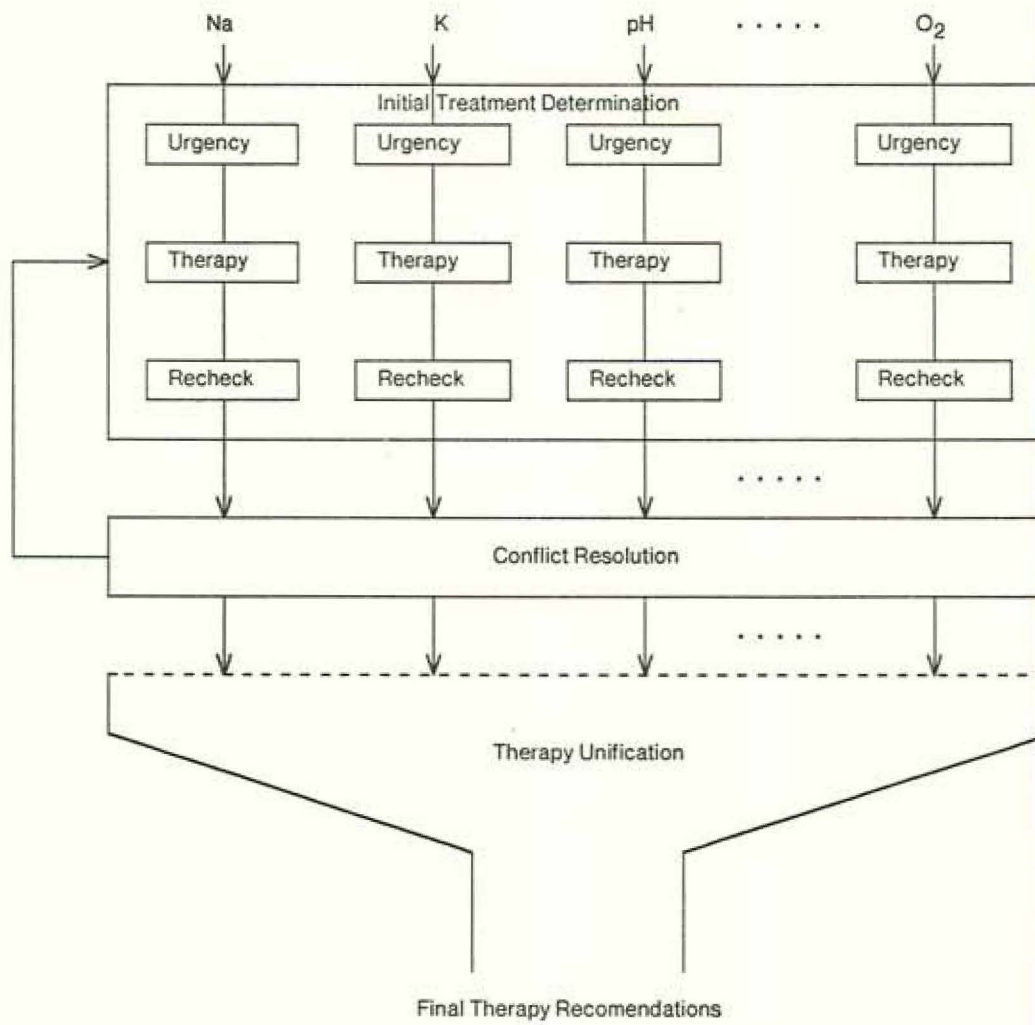
concentration of H^+ (rise in pH), and the concomitant drop in intracellular fluid H^+ , causes a shift of K^+ into the cells to avoid a net negative charge. Whenever H^+ moves across the cell membranes, K^+ moves across in the opposite direction in order to prevent a charge imbalance. Since the body's potassium stores are concentrated in the intracellular fluid, with much lower levels present in the bloodstream, movement of a small portion of the total body potassium in or out of the cells can have a dramatic effect on serum-k. So any change in serum-pH is accompanied by a significant change in serum-k: at normal levels, every increase of 0.1 in pH (reduction of H^+ by a factor of 1.25) causes a reduction of approximately 0.6 mEq/L in serum-k as H^+ leaves the cells and K^+ is pulled in to replace it.

Should abnormal distributions occur for any of these electrolytes, a full correction, while an appropriate long-term objective, is generally not an appropriate short-term goal. This is because the body will have partially adjusted to the imbalance and a rapid return to normal levels may result in adverse effects due to the excessive strain it would impose on the body's regulatory mechanisms. The short-term objective is therefore almost always to bring the levels back to normal as gradually as is consistent with the need to keep the patient out of immediate danger, allowing the body's homeostatic mechanism to adjust as changes occur, and to take over as soon as it safely can. Our formalism is well-suited to this desideratum, as the constraint of reducing the urgency one step at a time ensures that any interventions recommended would be only as abrupt as they needed to be to improve the patient's condition.

3. The Algorithm

The formalism presented above for formulation of a symptomatic therapy recommendation can be divided into three steps: (a) making independent recommendations for each abnormal condition, (b) resolving any conflicts due to interactions among these recommendations by revising some or all of them, and finally (c) combining all the individual therapy recommendations into a comprehensive treatment plan including the modality of treatment, the starting time and a time for reevaluation of the treatment and possible correction. (Figure 4 illustrates the three steps.) For our purposes, "abnormal condition" is taken to refer to an imbalance in a single electrolyte, so the individual recommendations are each concerned with one electrolyte.

Fig. 4. A schematic for overall system



3.1 Initial Recommendations for Individual Abnormalities

Arriving at a recommendation for an individual abnormality (step (a) above) is itself a three-step process: determination of urgency, initial treatment selection, and determination of recheck time. Each of the three steps is described below. All three involve some sort of initial estimate based on a crude characterization of the patient's condition, which is then revised to take into account the specific features of the case. This division into subtasks allows corresponding division of the production-rule knowledge base (see *Implementation* chapter for details) into largely independent segments, with well-defined paths of communication. We have, for instance, a clearly delineated body of knowledge related to determining the urgency of potassium disorders, another for selecting a specific treatment for such disorders, and a third for determining the recheck times. Three more distinct bodies of knowledge are related to performing the same tasks for pH disorders, and so on.

3.1.1 Determination of Urgency

The urgency of the need for therapeutic intervention provides the context for all later considerations. Based on our observations of clinicians, we chose to have five levels of urgency, ranging from stable (rating of 0) to critical (rating of 4). A patient is said to be stable when it is within the power of his/her body to control the situation, although some mild assistance may be in order. A critical patient is in immediate danger and requires substantial assistance without delay.

The determination of urgency consists of two steps, approximation and refinement. An initial approximation of the urgency for treatment of a given symptom is obtained solely on the basis of its severity. In our case, this refers to the concentration in the bloodstream of the electrolyte under consideration. This estimate is then adjusted, through application of the production rules, to reflect the specific circumstances of the case. These rules represent the heuristics mentioned above for assigning regions of the patient state space to patterns of clinical conditions. For example, a patient with moderate hypokalemia (low serum potassium) may start with a low urgency rating, but if the patient is also under treatment with a digitalis preparation, the rating will be revised substantially upward, because of the greater susceptibility of a digitalis-treated heart to hypokalemia-induced arrhythmias. After having been adjusted according to the expertise of the knowledge base specializing in this task, the urgency rating is finalized and made available for its role in treatment selection.

3.1.2 Determination of Initial Treatment Recommendation

The first step in treatment determination is the selection, from the program's repertoire, of those treatments which are applicable to the current case. These *active treatments* are chosen on the basis of the diagnosis and the urgency rating. (Again, the diagnosis and urgency rating are with respect to a single electrolyte, as recommendations are made for each electrolyte independently before being combined.) As it becomes active, each treatment is given initial ratings for risk and benefit. These ratings are also dependent on the diagnosis and urgency; a given treatment will have different initial risk/benefit ratings for the different contexts in which it can be activated. The urgency thus affects treatment selection both through the choice of active treatments and through their initial risk/benefit ratings.

Once activated and initialized, the plausible treatments have their ratings adjusted to the characteristics of the particular case. Much as in the determination of urgency, the adjustment is accomplished through application of a set of production rules specializing in the task. For those electrolytes whose treatments are almost completely determined by the diagnosis and urgency, the changes to initial ratings are minor. For others, with more variation in treatment according to the specific clinical context, the changes may be extensive. Potassium is an example of the latter, with many factors involved in the choice of treatment, and consequently the potential for pronounced departure from the initial ratings. With pH, on the other hand, the decision to treat is based almost exclusively on the extent of the imbalance, and once made, can be implemented in very few ways; the initial ratings according to urgency and diagnosis are therefore generally adequate, and undergo little change during rule application.

After all the rules have completed their adjustments, the active treatments are rank-ordered by the difference between their benefit and risk ratings. The treatment with the highest score (largest excess of benefit over risk) is then recommended as the treatment of choice. Should there be a tie, treatments with less risk are preferred. If the second highest ranking treatment has a score comparable to that of the first choice, that treatment is also recommended as an alternate. This safeguards against the imprecision of the scoring mechanism; the clinically appropriate treatment would have to be grossly misscored to miss being brought to the attention of the physician at least as an alternate.

3.1.3 Determination of Recheck Time

An initial estimate of the recheck time is based on the choice of treatment and on how quickly measurable effects are likely to appear. The more quickly change is expected, the sooner a reevaluation is in order. As above, the initial estimate is modified to fit the details of the current case via application of a set of rules designed specifically for that purpose. The rules are concerned with factors that would increase the dangers of allowing the patient's state to develop unmonitored. These considerations would include the severity of the condition, and the magnitude of the possible ill effects of overshooting or undershooting in the attempt to return to normal levels. The suggested recheck time, added to the chosen form of treatment, completes the set of recommendations for a single electrolyte.

3.2 Resolution of Cross-treatment Conflicts

Any inherent conflicts among the individual treatment recommendations must be resolved before they can be combined into a single comprehensive recommendation package. This process is initiated by the individual treatment formulation modules for each electrolyte, which may have *concerns* about any of the recommendations for other electrolytes. A concern is a list of conditions which, if present, would interfere with the therapy planned by the module having the concern, *i.e.*, would constitute a conflict. After all the independent treatment modules have run, the conditions are evaluated; those which are violated by recommendations for other electrolytes are activated as *objections*.

Objections are settled by modifying the recommendations which are in conflict. The module which raised the objection (as a concern) and the one(s) which violated it are identified and requested to revise their recommendations to resolve the conflict. Upon receiving such a request, a module will most often change its recommendation. If, however, the requested modification would adversely affect its attempt to regulate a critical abnormality, the module may refuse. If none of the conflicting modules are able to accommodate, the objection remains outstanding. All outstanding objections are reported to the user as unresolved conflicts. In such a case, the physician has to choose from among a set of alternatives none of which are entirely satisfactory. S/he must decide which aspect of the patient's condition is to have priority.

It is hoped that this method of resolving conflicts among independent recommendations may be applicable to other branches of medicine, as well. It commonly occurs that measures to correct different aspects of a patient's condition interfere with each other. If more than one requires immediate attention, some sort of compromise is necessary, much as in the treatment of electrolyte disorders. This difficulty can arise, for example, during the use of a mechanical ventilator. The ventilator provides breaths with pre-specified oxygen content, at a pre-specified volume and rate, to patients with poorly functioning respiratory systems. The volume and rate must not be set too high, or the pressure will interfere with blood flow into the heart and perhaps even rupture the lungs. Maintaining a sufficient supply of oxygen may therefore require a fairly high concentration in what volume is allowable, but too high an oxygen concentration is toxic. Meanwhile the pH, which is influenced by the rate of elimination of CO₂, must also be kept within narrow bounds. One way to resolve these complications might be to have separate modules responsible for the pressure, O₂ concentration, and pH. Each could recommend ventilator settings that would be satisfactory with respect to its own parameter, and raise objections to the settings favored by the other modules. As long as each could predict the impact of any particular ventilator setting on its own parameter, and could determine to what extent the patient's stability depended on the value of that parameter, they could conduct the same sort of negotiation used in the present program.

3.3 Final Unification of Recommendations

Finally, after conflicts between treatments have been resolved to whatever extent is possible all the individual recommendations are combined into a single comprehensive therapy package. This package designates specific therapeutic interventions as well as when and what follow-up examinations to perform in preparation for reevaluation.

3.3.1 Unification of Therapy Modalities

The combining of therapy modalities involves considerations of convenience, chemical interactions, and availability of standard solutions. The main factor in convenience is minimization of the number of distinct routes and times of intervention. Existing IV lines are taken advantage of, and any new IV solutions are combined as much as possible. The existence of IV line(s) may also be reason for switching an oral therapy to an intravenous additive.

Chemical interactions play a limiting role in the combination of solutions. Care must be taken to avoid possible interactions among the different preparations. For example NaHCO_3 and CaCl may precipitate if given simultaneously.

An attempt is made to use standard solutions because special formulation of solutions not routinely available introduces the possibility of error, and also means a potentially harmful delay.

3.3.2 Unification of Recheck Times

The second half of the comprehensive package is the follow-up recommendation. The independent recheck requirements for each electrolyte are translated into a series of bedside examinations and laboratory tests designed to gather all the information that will be needed to reevaluate the treatment. Convenience considerations similar to those mentioned for choosing therapy modalities are also involved here. It is desirable, from both the patient's and the physician's points of view, to minimize the taking of blood samples and other forms of examination. Consequently the measurements for each electrolyte are combined into a minimal number of lab tests selected from the standard repertoire, subject to the timing constraints imposed by the different recheck requirements.

The final suggestion for when to reconsult the program for evaluation of the patient's progress and the appropriateness of the treatment is taken to be the time of the earliest lab test, as this is based on the recheck requirements of the electrolyte(s) being measured. Once the unification is completed, the recommendations are presented to the physician.

4. The Implementation

ABET has been implemented in Maclisp [7] (using the LSB [1] and BrandX [14] extensions to the language) at the MIT Laboratory for Computer Science. The program's knowledge of medicine is encoded in production rules. But unlike other rule-based expert systems [2, 6, 12] ABET does not use goal-directed backward- or forward-chaining to identify applicable rule sets. Instead, as discussed above (section 3.1), we have divided the overall task into a number of relatively simple and isolated subtasks, and have relied on this partitioning of the problem to select the appropriate rules. The rules dealing with each subtask are packaged together, and the rules within any package may be executed in any order. When a package of rules is called to perform its subtask, each rule in it is tried exactly once. (The entire

package may be called again if the electrolyte it deals with is involved in a conflict.) To our surprise, no further organization of rules was found necessary. Should further development prove this structure to be inadequate, however, the program could easily be modified to introduce new rule-selection techniques.

There are three kinds of rules, one each for the three functions of the independent electrolyte-specific modules: determining the urgency, the initial treatment recommendations, and the recheck times. Their formats are similar. Each rule consists of an identifier (id), a list of predicates, a list of actions, and optionally a message. The id, an integer, is used for keeping a record of which rules have fired. The predicate portion has the same format for each kind of rule, and is described in the next section, on patient-specific knowledge. The actions are different for each kind of rule, and are described in the subsequent section. The message is simply a text string which is printed if the rule fires, mostly used to alert the physician to the presence of a condition which calls for some precaution.

4.1 Representation of, and Access to, Patient-Specific Knowledge

Eventually, as the complete patient management system for acid-base and electrolyte disorders takes shape, the therapy component (ABET), the diagnostic component (ABEL), and the top level decision module will all use the centralized *Patient-Specific Model* to represent what is known about the state of the patient. This multi-level, causal description of the patient's condition, discussed in [8], was designed for exactly such uses. But for ease of development, ABET has been given its own temporary - and much simpler - representation of the patient state. This allows us to defer the difficulties of working out an operational interface, and instead concentrate on the more immediate problems intrinsic to the therapy component itself. More details on how we expect to use ABEL's Patient-Specific Model can be found in section 6.3 (Suggestions for Further Work).

We simply consider the state of the patient to be summarized by a collection of descriptive characteristics. Each characteristic is associated with a value, which may be a number, an English word or phrase, or true or false. The complete catalog of patient characteristics is shown in figure 5. No structure is imposed on this set - there are no explicit relationships between characteristics.¹ The only operations

1. There are implicit relationships in that the values of some characteristics are calculated from others (see below), and are thus constrained to meet some mathematical relationship. But beyond the original derivation of the value, these links are not present or exploited in any explicit manner.

involving these characteristics are that each may be given a value (typically through the user answering a question), and each may be asked its value (typically for evaluation of a rule's predicate).

A query for the value of some characteristic will be handled in one of three ways. If the value is already known, it is simply returned. If it is not known but can be inferred from other characteristics, queries are generated for their values, the inference is made from the answers, and the inferred value is both returned and stored for later reference so the inference will not need to be repeated. The third case

Fig. 5. The patient characteristics

<u>Characteristic</u>	<u>Values</u>	<u>Source</u>
serum-k	<a number>	asked
serum-Na	<a number>	asked
pH	<a number>	asked
serum-glucose	<a number>	asked
K_diagnosis	low_K, normal_K, high_K	inferred from serum-k and dig_therapy
Na_diagnosis	low_Na, normal_Na, high_Na	inferred from serum-Na
pH_diagnosis	normal pH, respiratory alkalosis, metabolic alkalosis - chloride responsive, metabolic alkalosis - chloride unresponsive, respiratory acidosis, metabolic acidosis - acute, metabolic acidosis - chronic, metabolic acidosis - RTA type I, metabolic acidosis - RTA type II, metabolic acidosis - RTA type IV	asked
K_tempchar	acute, chronic, neither	asked
Na_tempchar	acute, chronic, neither	asked
pH_tempchar	acute, chronic, neither	asked
EKG	normal, abnormal, widened-QRS, peaked-T-waves, U-waves	asked
renal_function	none, mild, moderate, severe	asked*
dig_therapy	<yes or no>	asked
IV	<yes or no>	asked
npo**	<yes or no>	asked
hospitalized	<yes or no>	asked
seizures	<yes or no>	asked
respirator	<yes or no>	asked
body_weight	<a number>	asked
CHF	none, mild, moderate, severe	asked
fluid_outs	<a number>	asked
dehydration	none, mild, moderate, severe	asked
K_excess	<a number>	inferred from serum-k, pH and body_weight
ECF_deficit	<a number>	inferred from serum_Na and weight

*in the case of renal function, what is actually asked for is the creatinine clearance, from which one of the listed words is chosen to describe the degree of renal failure

***non per orum*, unable to take medication by mouth

is when a value is not known and cannot be inferred. Then the user is asked to supply the information. It is returned and stored for later reference.

When asking questions of physicians, it is important to conform to common medical interrogative style; should the program appear to be presenting questions in a random or incohesive manner, the physician will assume it has no cohesive idea of what it's looking for. Eventually, when ABET is connected to the diagnostic component (ABEL), questions will be directed through ABEL's information acquisition module, which will collect the inquiries and ensure that they appear in an acceptable order. Currently questions are just asked as they come up, but it should be noted that they still do tend to appear in a coherent and medically appropriate order as a result of the organization of rules into separate and conceptually cohesive groups.

Most queries occur in the course of the evaluation of the predicates in rules. The predicate portion of each rule is a list of one or more individual predicates, all of which must be true for the rule to fire (*i.e.*, for the action specified in the rule to be carried out). A single predicate is a triplet, consisting of a relation, the object being tested (the patient description or a particular aspect of it), and a value [16]. A predicate may be negated by preceding it with "not." If the relation is one of the characteristics used to represent the patient state, the meaning of the predicate is "Does the object have a value for this characteristic equal to that specified as the third part of the triplet?" An example of this kind of predicate is `[dehydration patient_state severe]`, which asks whether the "dehydration" characteristic for the patient has a value of "severe."

The relation may also be "greater" or "less," appearing with an object referring to a single characteristic of a patient description, as in `[less [serum-k patient_state] 3.0]`. The remaining relation is "present." Predicates using the "present" relation have as their third part the name of a characteristic whose only possible values are *true* and *false*. For one of these characteristics, rather than asking the awkward question of whether its value is equal to "true," one simply asks if the characteristic is present: `[present patient_state hospitalized]`.

4.2 Actions of Rules

The differences in the actions of the three kinds of rules reflect the differences in the data structures used to represent the urgency, the treatment choices, and the recheck recommendation. All three involve some sort of scoring. With each, some initial rating is adjusted by whichever rules apply to the case. A *history* records all changes made to the ratings - whenever a rule modifies a rating, its id number is added to a list of rules which have changed that rating.

This history is specifically designed to allow us, in the future, to provide explanations of the program's decisions. If every rule had associated with it a sentence or two giving the rationale behind its actions, then a justification of the ratings' initial values followed by the rationales for each of the rules that had changed them would constitute a fairly complete account of the process by which the program had arrived at its recommendations. All that would be missing is an explanation of the decisions made by the objection handler and the funnel.

The following sections describe the actions of the three kinds of rules and the data structures they affect.

4.2.1 Urgency Rules

In the determination of urgency, the urgencies corresponding to all possible values of the electrolyte under consideration are calculated simultaneously. The relevant data structure is the *urgency vector* (see figure 6). Each element of the urgency vector holds the urgency rating for a pre-defined range of electrolyte serum levels, with the complete vector covering the entire range of possible values. With potassium, for instance, the lowest range is for serum-k less than 2.0 mEq/L, the next is for serum-k between 2.0 and 2.4, and so forth to the highest, which is for serum-k greater than 6.75. All the elements are initialized to pre-set values, with high urgencies at the extremes and low ones in the center. The action of an urgency rule is to modify any or all of the values in the vector. A wide variety of operations to do so were initially implemented. We have found, however, that the only ones we still use are

subtraction, the setting of floors, and keeping of the old value (the null operation).²

A typical potassium rule and its actions are shown in figure 6. As can be seen from its predicate, this rule is invoked if the patient is taking a digitalis preparation. Its actions are to set floors on the urgency ratings corresponding to low levels of potassium and to leave unchanged those corresponding to high levels. This rule embodies the medical fact that digitalis treatment increases the heart's susceptibility to hypokalemia-induced arrhythmias. In one sense, this rule can be said to encompass three distinct rules about changes to urgency ratings for three different ranges of serum-k (under 2.4, 2.4 - 3.5, 3.5 - 4.0). Our use of the urgency vector to calculate urgencies for all ranges simultaneously thus allows us to combine several similar rules into one, making the knowledge base more compact and more understandable.

Once all the rules have been tried, the element of the urgency vector corresponding to the current serum level of the electrolyte under consideration is taken to be the urgency for treatment with respect to that electrolyte.

4.2.2 Treatment Rules

In treatment selection, on the basis of the diagnosis and the previously determined urgency, a set of specific interventions is activated. Each treatment has its own name, and as each is activated it is given initial risk and benefit ratings, also on the basis of the diagnosis and urgency. The treatment rules adjust these ratings. The actions listed in a treatment rule designate which treatment's ratings to modify, so unlike the format of urgency rules, in which null operations must be retained as place-keepers, only those treatments which are affected must be included, and they may be listed in any order. Each action consists of the name of the treatment and instructions for modifying the risk rating and the benefit rating. Many kinds of modifications were considered, but we again found it sufficient to allow subtraction and the setting of floors. If the treatment is not active, nothing is done; if it is active, its ratings are updated as

2. The array actually used to represent the urgency vector has two entries for each one of the ranges described in the text. This is to prevent ambiguities in the rating for a range which has been subtracted from by one rule and had a floor set on it by another. The convention we use is that regardless of the real chronological order of the rule applications, the rating for such a range is as though all of the settings of floors had occurred *before* any of the subtractions. One entry is used to keep track of each kind of operation: one for the strictest (highest) floor and one for the total of all subtractions. We nonetheless speak of a single urgency rating for this range because any time its rating is requested, the offset is automatically subtracted from the floor. Similar hidden mechanisms are used in scoring for the treatment and recheck recommendations.

Fig. 6. An urgency rule and its application

```
(add_urgency_rule '([present patient_state dig_therapy])
                  '((atleast 4) (atleast 4) (atleast 3) (atleast 3)
                    (atleast 2) (keep) (keep) (keep) (keep) (keep)))
```

Range	Serum-K	Urgency vector (initial)	Rule applied (dig therapy present)	Urgency vector (after rule application)
0	< 2.0	4	≥ 4	4
1	< 2.4	3	≥ 4	4
2	< 3.0	2	≥ 3	3
3	< 3.5	1	≥ 3	3
4	< 4.0	0	≥ 2	2
5	< 5.5	0	—	0
6	< 5.85	1	—	1
7	< 6.45	2	—	2
8	< 6.75	3	—	3
9	≥ 6.75	4	—	4

indicated in the action. For notational compactness, we use a positive number for setting a floor, a negative number for subtracting, and a 0 for specifying no change.

A typical treatment rule and its effects are shown in figure 7. This rule, from the sodium module, fires if the serum-Na is 125 mEq/L or less and the patient is suffering from seizures. As this would be an extremely critical situation, the rule's action is to favor the most aggressive intervention for low sodium. It sets the benefit rating for 3% NaCl IV ("Na_{low-rapid IV}") to at least 5, and makes sure this treatment is known to be slightly risky by setting a floor of 1 on its risk rating. Meanwhile the rule subtracts 4 from the benefit rating for the second most aggressive treatment for hyponatremia, "Na_{low-slow IV}," and leaves its risk rating unchanged.

As described in section 3.1.2, the active treatment with the greatest difference between its benefit and risk ratings after all the rules have been tried is considered the treatment of choice, and if the next best treatment has a score comparable to that of the first choice, it is recommended as an alternate.

Fig. 7. A treatment rule and its application

```
(add_treatment_rule '([less [serum-Na patient_state] 126]
                    [present patient_state seizures])
                    '((|Na,low-rapid IV| 1 5) (|Na,low-slow IV| 0 -4)))
```

Treatment name	Initial scores		Rule applied		after rule application	
	risk	benefit	risk	benefit	risk	benefit
Na, none						
Na, low-slow oral						
Na, low-mod oral						
Na, low-rapid oral*	0	0			0	0
Na, low-slow IV*	0	0	0	-4	0	-4
Na, low-rapid IV*	0	0	1	5	1	5
Na, NS*	0	0			0	0
Na, D5half						
Na, half NS						
Na, high-oral						
Na, high-IV						

At urgency = 3 (very unstable), and with a diagnosis of hyponatremia, the starred (*) treatments are active.

4.2.3 Recheck Rules

The format of the actions for recheck rules is simpler since only one variable is affected. The recommended recheck time will be a value between 0 and 6. Zero indicates a recommendation to reconsult in an hour or two; five indicates a waiting time of one to two weeks, and six indicates no follow-up is needed (for the electrolyte under consideration). The intermediate steps vary by roughly two-fold per step. The initial value for recheck time is based on how quickly the treatment of choice is expected to produce noticeable effects. This estimate is then adjusted by application of the recheck rules, which are concerned with the importance of monitoring the development of the patient. Each recheck rule may add to or subtract from the existing value, or set a ceiling. A typical recheck rule can be seen in figure 8. This rule states that if the patient's pH is less than 7.2, a ceiling should be placed on the recheck recommendation of one to two hours (recheck time "0"), and a message printed for the user, reminding him/her of the severity of the situation.

Fig. 8. A recheck rule

```
(add_recheck_rule '([less [pH patient_state] 7.2])
                  '(nolongerthan 0)
                  '|Severe acidosis requires recheck every hour|)
```

4.3 Conflict Resolution (The "Objection Handler")

As discussed in section 3.2, conflict resolution begins with the raising of concerns by the independent modules responsible for each electrolyte. A concern is raised when the potential exists for another module to recommend treatment which would interfere with that planned by the module having the concern. The modules each have a *concern generator* which runs after the module has reached its own initial recommendation. The concerns it raises depend on the patient's condition and on the present module's recommendation. As the other recommendations are not yet necessarily available, it cannot be known whether the predicted interference will actually arise. Consequently a concern must include some test, to be performed after all the initial recommendations have been made, for recognizing whether the potential conflict actually arose as foreseen. A concern consists of (a) a name, to identify the conflict it addresses, (b) a description of the conditions under which the potential interference would arise, and (c) a list of which modules can resolve the conflict by revising their recommendations, typically the module having the concern and the one responsible for the interference.

After all the independent modules have given their initial recommendations, the *objection handler* considers whatever concerns have been raised. Each one whose test conditions (part (b), above) are present becomes an active objection.³ The objection handler then asks each module in the list of disputants (part (c), above) how willing it would be to settle the objection by revising its recommendation. The module responds with a number from 0 to 4, with 4 indicating absolute refusal. If every module asked refuses absolutely, the conflict is considered irreconcilable and is reported to the user, who then must decide which module (which aspect of the patient's condition) to favor. Otherwise the module which responded as being most willing is instructed to carry out the necessary modifications to its

3. This method for resolving conflicts is similar to the "goal protection" scheme used by Sacerdoti [10] to prevent the planning process for achieving one goal from interfering with the planning for another goal.

recommendation.⁴ Apart from raising its own concerns, each module is thus equipped to decide on the basis of the current clinical situation how willing it is to change its recommendation to resolve some conflict, as well as actually proceed with the modifications.

To clarify the operations involved in conflict resolution, let us consider one example, the interaction between serum-k and pH. Recall that an increase in pH causes a substantial drop in serum-k, and a decrease in pH causes serum-k to rise.

In its initial recommendations, the potassium module assumes a constant serum-pH. One of its concerns will be that if serum-pH is abnormal, the pH module may recommend treatment to change it. When the objection handler considers this concern, it checks the pH recommendation; if the pH can be expected to change because of the recommended intervention (or for some user-supplied reason outside the program's domain), the concern is made into an active objection. The potassium and pH modules are identified in the concern as capable of settling the objection by revising their recommendations, so each is then asked how willing it would be to do so.

The response from the pH module is simply the previously determined urgency for pH. (There is one exception: acidotic conditions with urgency 3 respond with a 4.) So the less stable the patient, the less willing the pH module is to accommodate. The response from the potassium module depends on the direction of the expected change. If the shift would tend to alleviate the potassium problem (bring serum-k closer to normal), the answer is 0; the K-module would be more than happy to have an assist. If the shift would push serum-k further from normal, the response is the urgency rating for the predicted serum-k, which will be greater than or equal to the current rating.

If both responses are 4, the conflict is reported to the user as unresolved. If not, the more willing module is instructed to modify its recommendation.⁵ Should this be the pH module, it will cut back its recommended intervention to prevent aggravating the potassium disorder. If it is the potassium module, it will in most cases prescribe a new treatment based on the predicted blood chemistry. The exception is

4. In case of a tie, the module which comes first on the list is chosen.

5. Currently, one module is chosen and instructed to change its recommendation enough to resolve the conflict. Although it hasn't yet become necessary, we expect we will eventually need the capability to have several modules each modify their recommendations less drastically, and thus resolve the conflict by cooperatively meeting partway.

that when the patient is currently critical (urgency = 4), it will not lessen the aggressiveness of its recommendation even if the prediction is for an improved serum-k; we cannot afford to wait.

4.4 Final Unification (The "Funnel")

There are two parts to the final combining of the independent recommendations: unification of therapy modalities, and unification of follow-up recommendations. The emphasis in unification of therapy modalities is on combining multiple IV solutions into a single carrier with additives. By carriers we mean solutions of NaCl and/or dextrose with tonicity close to that of the bloodstream. Administration of such solutions has little effect on blood chemistry, so dissolving small quantities of more potent substances in carriers is a good way to ensure safe concentrations of the additives.

The first step is to identify a carrier. The treatments for pH or for sodium and water balance may already involve a suitable solution, and if so we just use that. Otherwise we choose a carrier. The default is "DSW," an isotonic 5% (by weight) solution of dextrose. But if blood sugar is already high, we instead use "½NS" ("half-normal-saline"), a NaCl solution with tonicity half that of the bloodstream.

Once settled on a carrier, the additives required by the various treatment recommendations must be considered. Some can simply be mixed in (*e.g.*, moderate amounts of potassium chloride for hypokalemia); some require their own line (*e.g.*, isotonic hydrochloric acid for severe metabolic alkalosis); a few require special consideration. Bicarbonate supplementation, to alleviate acidosis, is in the form of NaHCO_3 . This is added to the carrier, but the presence of the accompanying sodium may necessitate switching to a carrier with lower sodium content. And when large amounts of potassium chloride are needed immediately, mixing it with the carrier may be too slow - a special arrangement known as a "piggyback" is used instead.⁶

6. Piggybacking refers to an infusion through an established intravenous catheter connected to another bottle. A Y-shaped piece of tubing is inserted into the main tubing; the piggybacked solution is run in one branch of the "Y" and the main solution into the other. The mixture flows out the bottom of the "Y" and into the patient. Piggybacks are used when many different solutions have to be infused over a short period of time, when urgent treatments must preempt ongoing less urgent ones, or when highly concentrated solutions must be used.

Unification of follow-up recommendations involves choosing a battery of bedside and laboratory tests to gather the information that will be needed to reevaluate the treatment when the program is reconsulted. Selection of bedside examinations is quite simple: each of the physical tests the program knows about is associated with a predicate. If the predicate is true, the test is suggested. Thus for dehydrated patients, it is suggested that urine output be monitored; patients suffering from acidosis should have their respiratory rate watched; *etc.*

Choosing laboratory tests is a matter of minimizing the taking of blood samples while ensuring the availability of timely data. Each of the independent modules will have decided on a recheck time for its electrolyte. Any two of these which are within one step of each other (recall that each step represents roughly a doubling of the waiting time) are combined into a single blood sampling. Then the actual tests to be performed are chosen. For pH, there are no decisions - an arterial blood gas is recommended at whatever time the information is needed. But for other serum electrolytes, any of several tests may be specified. They vary in how long they take, and in how broad a range of electrolytes are reported. If the potassium module had asked for a recheck in two hours, and no other electrolytes need be known about so soon, the recommended test is the Na/K determination, which doesn't take long but only reports on sodium and potassium. On the other hand, if the potassium module wanted a recheck in 24 hours, and at that time information on calcium would also be needed, the suggested test would be the chemprofile, which reports on six electrolytes and takes the lab about half a day to process.

5. An Example

The case description for our example is shown in figure 9. With respect to acid-base and electrolyte balance, the patient's most pressing problem is severe acidosis. The potassium level is slightly low, and can be expected to worsen considerably as the acidosis is corrected. Treatment is also necessary for moderate dehydration (and its attendant hypernatremia). Obviously the underlying diabetes and pneumonia require attention as well, but that is well beyond the domain of our program.

The example is presented and commented on step by step. Underlined text indicates user input in response to questions; anything not underlined is program output. As discussed above (in section 4.1), although we ultimately expect ABET's interrogation to work through an information acquisition module, for now it simply asks for information as it needs it.

Fig. 9. Case Description

A 65 year old man with a history of diabetes mellitus controlled with diet, hypertension treated with a diuretic, and mild congestive heart failure presents to the emergency ward with pneumonia complicated by diabetic ketoacidosis of two days' duration. Physical examination reveals moderate dehydration, rapid breathing, fever, and signs of pneumonia. X-ray examination of the chest and sputum examination are consistent with pneumonia. Urine and blood studies are consistent with diabetic ketoacidosis. He is admitted to the hospital for further treatment.

The first step is the formulation of independent recommendations for each electrolyte. Figure 10 shows the activity of the pH module. The low pH of 7.15 classifies the patient as "very unstable," an urgency rating of 3. Given that level of stability, and the diagnosis of *acute metabolic acidosis*, there is no choice but intravenous administration of bicarbonate, *i.e.*, the *active treatments* consist of only one possibility. This treatment, "pH,IV_HCO₃_fast," needs to have its rate of administration computed. To do so, the program first asks for the patient's pCO₂ and body weight, then calculates a serum-bicarbonate of 7 mEq/L, 17 short of the ideal 24 mEq/L. Multiplied into the bicarbonate space of distribution (approximately 40% of body weight), this implies a total deficit of 480 mEq. Our initial goal is to replace half of that, or 240 mEq. At 44 mEq per ampoule, 5½ ampoules would be needed, but three is the maximum we can safely add to a one liter solution. The recommendation is therefore to dissolve three ampoules of NaHCO₃ in a liter of D5W. The presence of mild congestive heart failure triggers a message to the user about the hazardous possibility of overloading the circulatory system. Finally, the minimal recheck time of one to two hours is recommended, due to the severe acidosis. The message to that effect in the "Comments" area is from the appropriate recheck rule.

Figure 11 shows the formulation of the sodium recommendation. The slightly high serum sodium of 148 causes an initial urgency rating of 1 ("mildly unstable"). The EKG information is needed to evaluate the predicate of a rule which sets most ranges to "critical" if the EKG is abnormal. (The condition of *sinus tachycardia*, while reflecting an elevated pulse rate, is considered normal for our purposes.) The next question is for a rule which subtracts one from the urgency ratings for several ranges if the disorder is chronic (and if the EKG is normal and there are no seizures). Neither of these rules applies in our case.

Fig. 10. Independent Recommendations: pH

----- Formulation of Independent Recommendations -----

*** considering electrolyte PH ***

A. Urgency Determination

pH: 7.15

Regarding electrolyte PH, the patient is classified as: very unstable (grade 3 of 4)

B. Treatment Selection

What is the pH diagnosis? metabolic acidosis - acute

recommended:

IV bicarbonate as 1.26% solution over 3-4 hours

pH, IV_HCO3_fast

benefit rating: 5.0

risk rating: 4.0

score: 1.0

pCO2: 20

body weight (Kg): 70

Calculated serum bicarbonate is 7.0

Base deficit is 17.0 mEq/liter.

Total base deficit is 480.0 mEq.

Initial dose is 3 ampoules of sodium bicarbonate

Administer as 7.5% solution over 5 minutes if the patient is coding.

Otherwise, add to one liter of D5W and administer over 3-4 hours.

Comments:

Monitor carefully for congestive heart failure from salt overload.

C. Recheck Recommendation

recommended recheck time: one to two hours

Comments:

Severe acidosis requires recheck every hour

The dehydration question illustrates one of the program's convenience features. Each characteristic has associated with it a list of possible values. If the answer to a question does not match any of these values, the list of choices is displayed and the question is asked again. Besides catching typing errors, this feature makes it easy for the user to get a list of acceptable responses - all s/he need do is enter something which won't match any of them, such as the question mark in the example.

Fig. 11. Independent Recommendations: sodium

*** considering electrolyte NA ***

A. Urgency Determination

Na: 148

EKG: sinus tachycardia

Is the sodium disorder "acute", "chronic" or "neither"? acute

Degree of dehydration: 2

Please enter one of the following: NONE, MILD, MODERATE, SEVERE

Degree of dehydration: moderate

Regarding electrolyte NA, the patient is classified as: moderately unstable (grade 2 of 4)

B. Treatment Selection

Is the patient unable to take medication orally? n

serum-glucose: 500

recommended:

0.5 normal saline - rate to be calculated

Na, half NS

benefit rating: 5.0

risk rating: 1.0

score: 4.0

Degree of congestive heart failure: mild

creatinine clearance: 100

Estimated degree of dehydration is MODERATE

CHF = MILD

renal function = NORMAL

IV rate is slowed by factor of 1.5

Calculated optimal IV rate is 400.0 cc/hour.

C. Recheck Recommendation

recommended recheck time: one to two hours

Comments:

Moderate dehydration reduces recheck time 50 per cent

mild CHF reduces recheck time 50 per cent

Moderate illness requires recheck in 8 hours

The eventual response of "moderate" dehydration triggers a rule which sets the urgency to at least 2. The moderate dehydration indicates that even though the serum sodium is high, the patient actually has a sodium deficit; there is little enough fluid for a less than normal amount of sodium to yield a higher than normal concentration.

No other rules are applicable, and the final urgency rating is 2.

For treatment selection, apart from what we already know, there are questions about serum glucose, and about whether the patient is able to take medication by mouth. Had the patient been unable to take medication orally, perhaps due to persistent vomiting, all therapies which depend on that route would have been strictly ruled out. The serum glucose influences the choice of carrier. The highest scoring of the active treatments is half normal saline. Before calculating its rate of administration, the program needs to know about the presence of renal failure (measured by creatinine clearance) or congestive heart failure. With moderate dehydration we would like to administer 500 cc/hour, but the presence of mild congestive heart failure forces us to slow down to 400.

The recheck recommendation is again the minimal one to two hours. As can be seen from the comments in the example, the urgency of 2 (moderate illness) sets a maximum of 8 hours; this figure is cut in half by the presence of mild congestive heart failure and again because of the moderate dehydration.

Formulation of the potassium recommendation is shown in figure 12. The serum-k of 3.0, being slightly low, gives us an initial urgency rating of 1. The urgency rules for potassium are concerned with heart function and whether the disorder is acute or chronic. Since the patient is not taking digitalis, and is already known to have a normal EKG, we are not as worried as we might be about hypokalemia-induced arrhythmias. On the other hand, we are more concerned than we would be if the condition were chronic, because in chronic situations the body generally has had a chance to partially compensate for an abnormal electrolyte level. The final urgency rating is just what we started with: mildly unstable.

For treatment of mild hypokalemia, there are several oral and intravenous methods of supplementation, with varying rates of effect. The fact that the patient does not currently have an IV gives a slight edge to the less intrusive oral therapies. (The funnel will realize that the sodium and pH modules are recommending intravenous treatments, so if the potassium module chose an oral therapy

Fig. 12. Independent Recommendations: potassium

*** considering electrolyte K ***

A. Urgency Determination

Is the patient on digitalis? n

Is the potassium disorder "acute", "chronic" or "neither"? acute

K: 3.0

Regarding electrolyte K, the patient is classified as: mildly unstable (grade 1 of 4)

B. Treatment Selection

Does the patient have an IV? n

Is the patient hospitalized? y

Estimated potassium excess = -630

recommended:

KCl tabs, 80 mEq po qD

K, low-rapid oral

benefit rating: 5.0

risk rating: 0.3

score: 4.7

also consider:

KCl, 40 mEq IV over 24 hours

K, low-slow IV

benefit rating: 5.0

risk rating: 1.0

score: 4.0

C. Recheck Recommendation

recommended recheck time: one to two weeks

here, the funnel would probably switch it to an IV additive anyway.) Were the patient suffering from renal failure, aggressive treatments would be very risky, but we already know his/her renal function is normal. The last consideration is the magnitude of the potassium deficit. Although the serum-k is not far from normal, the very low pH indicates a large actual shortage of potassium, favoring active

intervention.⁷ All these factors combined make the most aggressive of the oral therapies the treatment of choice, with the slowest of the IV therapies a close second.

The recheck recommendation is a rather leisurely one to two weeks, despite the administration of potassium supplements, because the urgency rating is still only a one, and no drastic changes are expected. This picture will be dramatically altered when the objection handler considers the effects of the recommended treatment for pH.

Figure 13 shows the efforts of the objection handler to resolve the conflict between the potassium and pH modules over the serum-pH. The potassium module, as a matter of routine, had raised a concern called "K_SHIFT_DUE_TO_PH," just in case the pH should change enough to shift significant amounts of potassium in or out of the cells. To evaluate this concern's activation conditions (a significant potassium shift), it was necessary to predict the pH. The program made its own prediction by guessing that the recommended pH treatment would bring the pH halfway back to the normal 7.4 from its current value of 7.15. The user concurred with this estimate, which was then used to predict a new serum-k of 2.22. This is indeed a significant shift, and the concern became an objection.

The objection handler then asked the pH and potassium modules how willing each would be to revise its recommendations. As discussed in section 4.3, the pH module, when it is facing an acidosis with an urgency of 3 or greater, responds with 4, meaning "absolutely not." The potassium module, facing a pH change that will worsen the serum-k situation (push it further from normal), answers with the urgency rating that would correspond to the predicted serum-k. In this case, that is a 3. Since the potassium response was lower than the pH response, the potassium module is directed to revise its recommendation in light of the fact that the pH recommendation will be allowed to stand.

7. This consideration of the pH ought not be confused with the kind of cross-module conflicts that are dealt with by the objection handler. There we are concerned with interactive effects caused by the intervention recommended by another module. Objections are raised to prevent changes from happening behind the current module's back, due to the actions of other modules. Here, on the other hand, we are concerned with the static masking effect of pH on the body's total stores of potassium. We are simply acquiring a more complete picture of the current potassium situation than is possible from consideration solely of the serum-k.

Fig. 13. Conflict Resolution

```
----- Objection Handler -----  
  
predicted pH (estimated at 7.28 based on pH treatment recommendation): 7.28  
concern K_SHIFT_DUE_TO_PH activated as an objection  
  
pH willingness to revise: 4  
K willingness to revise: 3  
  
Reconsidering K  
  
predicted K: 2.22  
new urgency rating: 3  
  
Worsening of serum-k is expected on the basis of the predicted pH  
change. Correction of K disorder should be delayed until the expected  
pH change occurs.  
  
B. Treatment Selection  
  
Estimated potassium excess = -630  
  
recommended:  
KCl, 10 mEq IV over one hour times three doses  
K, low-slow crash  
benefit rating: 5.0  
risk rating: 2.0  
score: 3.0  
  
Comments:  
*** WARNING: This level of hypokalemia may be life-threatening!  
The patient should have an electrocardiogram done now and serially  
to evaluate the degree of physiologic hypokalemia.  
  
C. Recheck Recommendation  
  
recommended recheck time: 4 to 6 hours
```

What follows is a rerun of the potassium module, assuming the predicted serum-k. This time a very aggressive IV therapy is chosen. With an urgency of 3 the recheck recommendation is dropped to 4 - 6 hours, and the user is cautioned not to start until the predicted changes actually occur, lest we pump large quantities of KCl into a patient whose serum-k is already near normal.

One step remains: the unification of recommendations. Figure 14 shows the funnel in operation. The pH recommendation is passed through unchanged. On consideration of the sodium recommendation (400 cc/hr of 1/2NS, to alleviate volume depletion and sodium deficit), it is discovered that the bicarbonate solution already planned for acidosis will also satisfy the sodium treatment

objectives. Its rate of ~300 cc/hr is close to the 400 cc/hr recommended for volume depletion, and the sodium in NaHCO_3 will sufficiently reduce the sodium deficit. Since the bicarbonate infusion adequately serves the treatment objectives of both the pH and sodium recommendations, the $\frac{1}{2}$ NS solution is simply dropped. And the potassium supplement, because it requires rapid administration of concentrated KCl, becomes a piggyback on the bicarbonate solution.

Next come the follow-up recommendations. The presence of dehydration, hypokalemia, and acidosis call for the bedside examinations shown. As for laboratory tests, the sodium and pH modules have asked for consultations in one to two hours, and the potassium module in four to six hours. Since the potassium request is only one step longer than the others, rather than take blood twice the program recommends checking all three in one to two hours. It suggests the Na/K determination for sodium and potassium, and an arterial blood gas for pH, all from one blood sample. That concludes the case.

Fig. 14. Unification of Recommendations

----- Final Unification -----

Combined treatment recommendations:

IV Solution #1:
NaHCO₃, 150 mEq (three ampoules in D5W, one liter) to run over 3-4 hours

IV Solution #2: (piggyback onto IV Line #1)
KCl, 10 mEq in 50 cc D5W, to run over one hour (three doses).

Recommended bedside examinations:

For dehydration: vital signs, lungs, heart, urine output
For hypokalemia: weakness, electrocardiogram
For acidosis: respiratory rate (watch for fatigue)

Recommended laboratory tests:

For potassium and sodium: Na/K determination in one to two hours
For pH: arterial blood gas in one to two hours

END-OF-CASE

6. Discussion and Conclusions

6.1 Survey of Therapy Programs

Earlier attempts to prescribe therapy with AI techniques have often been subsumed in more general consultation programs that also conducted diagnosis. Two such programs, described below, have been written for diagnosis and treatment of infectious disease (MYCIN [12]), and of eye diseases (CASNET/Glaucoma [15]). Two other projects have concentrated entirely on treatment: the Digitalis Therapy Advisor [3], for management of patients receiving digitalis preparations, and VM [2], for management of patients on ventilators.

6.1.1 Mycin

MYCIN specializes in diagnosis and treatment of bacterial infection of the bloodstream. It uses associative triples to represent the patient-specific knowledge, and production rules for its medical knowledge. Each rule consists of a set of premises and a conclusion. The interpretation is that if the premises are known to be true, the conclusion is taken to be partly *confirmed* with a numerical measure of certainty. The same rules also direct the flow of control through goal-directed backward-chaining. The highest level goal is to determine if the patient is suffering from a significant infection which should be treated, and if so, to select the appropriate therapy. This goal, like all others, is pursued by first retrieving all the rules which state the goal as their conclusion. Sequentially for each rule in this set, MYCIN then attempts to confirm the goal by evaluating the premises of the rule. If the validity of a premise is not already available in the data base, then determination of the premise's validity itself becomes a goal. The program recursively pursues this new goal in the same manner, applying another set of rules. The resulting tree structure of hierarchical goals eventually reaches the level of primitive facts. If trying all the applicable rules for a particular goal fails to either conclusively confirm or deny it, MYCIN asks the user for the clinical information that will establish the validity of the goal.

This backward-chaining goal structure allows efficient problem solving, but it also diminishes the modularity of the knowledge base, as it becomes necessary to consider the interactions between rules during problem solving. Rules must be written with an eye towards their role in dynamic construction of the goal structure as well as their role in representation of medical knowledge.

6.1.2 CASNET/Glaucoma

The Glaucoma program uses the CASNET theory of representation of causal knowledge to perform diagnosis and therapy of eye diseases. Medical knowledge is represented as a network of nodes corresponding to physiological states, linked together by transition probabilities. States are also associated with support values, indicating how strongly certain test results verify the presence of the state. The transition probabilities are used to decide which state would be most fruitful to investigate next, while the support values are used to score states for confirmation or denial. Diseases are recognized by comparing the pattern of confirmed and denied states to those of known disorders. The therapy alternatives are then evaluated through the same techniques used for diagnosis. A new causal network is constructed around nodes representing the various therapies, and their effects are probabilistically propagated through the network to determine which of the original disorders would remain.

CASNET's explicit reliance on causality is advantageous. People seem to prefer knowing why something happens to knowing merely that, under the circumstances, it does. A program that works in these familiar terms is more likely to appear sufficiently reasonable to be accepted by physicians. Moreover, perfecting the performance of such a program is a plausible goal, because any errors it makes in classifying a patient must ultimately be related to a flaw in the causal model. The mistake should be correctable by adding more tests to distinguish the erroneous case or disaggregating some state in the network into several to give a more detailed model of some aspect of the disease. In a more statistically oriented approach such local refinement would be difficult.

6.1.3 Digitalis Therapy Advisor

The drug digitalis is commonly used to strengthen and/or stabilize the heartbeat. Unfortunately, even slight overdoses have toxic effects. Since patient sensitivities vary, and the signs of toxicity are subtle, digitalis is difficult to administer safely. Toxicity occurs in 20 per cent of all patients receiving the drug [4]. Although experienced cardiologists often achieve a far lower incidence of toxicity, other physicians having less familiarity with the drug and the signs of its effects may encounter more episodes of toxicity. The Digitalis Therapy Advisor was developed with hopes of spreading more widely the knowledge of expert cardiologists regarding digitalis use, thus helping less experienced physicians use the drug effectively.

The program builds a patient-specific model, involving both a formal, quantitative pharmacokinetic model, and qualitative, clinical data. It then uses the patient model to formulate an initial dosage regimen, specifying a goal for total body stores, and loading and maintenance doses calculated to achieve and maintain that goal. In a continuing series of consultations with the physician, the program carefully monitors the patient's clinical responses for signs both of toxic and therapeutic effects. Its patient model and dosage recommendations are updated through a feedback process. The patient's actual response provides increasingly exact guidance as to his/her drug sensitivity and rate of absorption, allowing an increasingly accurate patient model and dosage recommendations.

The program's main improvements over earlier attempts in this field are twofold. First, in formulating its initial dosage recommendations, it takes into account patient-specific factors that might increase digitalis sensitivity, through the patient-specific model. Second, in the feedback process that refines its patient model, the program makes use of clinical as well as quantitative data. The nonquantitative clinical information, although harder for a program to codify and work with, if interpreted properly actually provides a more relevant portrayal of the patient's condition.

6.1.4 VM

VM is a program for the management of patients receiving ventilation assistance. Such patients need to be weaned off their ventilators, gradually increasing the demands made on their own respiratory systems. VM uses MYCIN-like production rules for both knowledge representation and control of program flow. The main difference is that VM uses forward-chaining instead of MYCIN's backward-chaining mechanism. That is, the rules to be tried next are selected on the basis of having premises which have just been confirmed, rather than on the basis of having goals which would confirm the premises of the rule which has just been tried. The strategy is "What can we prove with this?," rather than "How can we prove this?"

Like the Digitalis Therapy Advisor, VM models the development of the patient through time. The inclusion of time is its main improvement on the MYCIN-style rules. Premises often ask how long a certain condition has been met. But VM's model of development is limited by its low level of differentiation among patients. It attempts to guide each patient along a fixed path of improvement, which varies only according to the type of ventilation assistance and a handful of patient characteristics.

Without more flexibility, it is difficult to be prepared for the entire range of likely situations. These limitations are not, however, intrinsic to VM's fundamental approach - Fagan's thesis does mention possible extensions to provide more patient specificity.

6.2 Evaluation

What ABET embodies is the first general approach to therapy of acute symptomatic problems. In particular, this task has been accomplished for the domain of acid-base and electrolyte disorders via decomposition of the problem into a number of conceptually simple steps. ABET's performance, while adequate for an exploratory demonstration project, could certainly be improved through adoption of some of the techniques developed for the other therapy programs. Use of causality, whether in the form of the CASNET representation or in the form of ABEL's causal links, could lend our program a deeper understanding of its domain. Its knowledge is now largely phenomenological, without explicit reasons given for the correspondences between clinical conditions and preferences for certain urgency categories or specific treatments or recheck times. The rules expressing these correspondences are often identifiable with particular pieces of standard causal medical reasoning, but the causal reasoning itself is not directly exploited - it has been transformed into the rule set and is present only implicitly. Explicit inclusion of the medical reasoning behind ABET's actions would enhance the program's clarity and allow it to handle unanticipated interactions in a principled manner, as well as easing the development of explanation facilities.

Mycin's backward-chaining mechanism gives it an effective goal structure far more flexible than ABET's control mechanisms. We had originally anticipated needing something similar. However, as discussed earlier, the clinical practice of acute symptomatic therapy has a natural decomposition which allows us to perform the task without an elaborate scheme for sequencing of operations. Should our program prove unable to adapt to unforeseen clinical complexities, a new method for rule selection could easily be inserted.

A weakness of the program as it stands is the absence of quantitative reasoning. Some calculations are performed, but only in the peripheral matter of ascertaining the values of certain patient characteristics. No quantitative modeling is performed to predict the effects of therapies under consideration, nor to analyze the efficacy of previously administered therapies. Much could be learned,

in this respect, from the Digitalis Therapy Advisor and its pharmacokinetic modeling.

And ABET needs a conception of time something like those of the Digitalis Therapy Advisor and VM. In the original proposal for the Acid-Base and Electrolyte Consultant System [13] initial therapy formulation is viewed as only the first of three steps in the treatment procedure. It must be followed by an evaluation of the patient's response and an improved reformulation of the original therapeutic regimen. Before it can be an effective part of an overall patient management system, the therapy component will have to understand the patient's progress through time, perceive the inadequacies of its own recommendations, and know how to adjust. The Digitalis Therapy Advisor, in particular, has been very successful at tracking the patient through time, because its patient model is flexible enough to represent the subtleties of a specific patient's development.

In one sense, however, these criticisms are irrelevant. The completion of a system capable of expert-level performance was not the intent of this project. We wished to explore the feasibility of applying AI techniques, and particularly a certain domain-independent model of the therapy process (that outlined in the *Task Domain* chapter), to the field of acid-base and electrolyte balance. What we have found is promising. The program itself has achieved a basic level of competence, using simple and hopefully generalizable methodology. The field of acid-base and electrolyte balance has been prepared for further work, and our model of the therapy process has shown itself applicable to at least one domain. The fact that other projects, aimed specifically at high-level performance in a single field of medicine, have been successful with techniques we have failed to use is an indication of how we might improve ABET's performance. It does not, however, bear directly on the successfulness of our own project.

6.3 Suggestions for Further Work

The current program's continuing usefulness as an exploratory research tool requires a variety of improvements, some minor and some rather ambitious. A few of these have been discussed above, in the course of comparing ABET to other therapy projects: use of causal and quantitative models, adoption of a more flexible goal structure, and incorporation of some sense of the patient's development through time. There are also a number of other possibilities.

A few direct extensions to what's been done, simple but nonetheless likely to yield interesting results, would be the addition of more electrolytes to ABET's repertoire, and trial of the objection handler on more extensive interactive effects. These would allow a more thorough trial of the program's performance in its target domain. Other extensions, requiring new facilities but little change to the existing program, would include an explanation generator and an information acquisition module. These improvements, although having little impact on ABET's performance capabilities *per se*, would be essential prerequisites to the program's acceptance in the clinical setting. Physicians would refuse, quite reasonably, to heed the advice of a program which could not explain its reasoning in their terms, or which did not appear to go about the task of information gathering in a familiar or conventional style.

Modifications of a more fundamental nature are needed to help ABET more closely simulate the cognitive processes of actual clinical practice. Some sort of verification or consistency-check ought to be done on symptomatic findings. The program ought to notice and question further when conditions which don't generally coincide seem to be simultaneously present in the description of the patient state. To recognize this sort of situation reliably may well require a deeper functional understanding of physiological mechanisms than the program now has. It is perhaps possible that the deeper knowledge embedded in the diagnostic component, ABEL, could be utilized to determine the plausibility of the patient state presented to ABET.

It will also be necessary to extend ABET's expertise to etiological treatments. We currently envision doing so by having ABEL provide a simple causal network describing the mechanism underlying the patient's disorder, distilled from ABEL's more elaborate Patient-Specific Model [8]. We would then step through the network and select the earliest point at which we could intervene, whether it be a node representing a dysfunctional state or a link connecting two such states.⁸ After choosing an appropriate means of intervention, we predict how the acute situation will evolve in light of our treatment, and apply the symptomatic portion of ABET to the difficulties that are expected to remain.

8. A similar method for isolating and treating the cause of a disorder is used in the Congestive Heart Failure program [5].

Many of these improvements will occur as a consequence of coupling ABET to the rest of the Acid-Base and Electrolyte Consultant System. We could then employ the facilities already present in the ABEL diagnostic component. Its information acquisition module could ensure a medically appropriate style for our questioning. Use of its quantitative modeling capabilities could improve ABET's understanding of the effects of therapies, both those projected and those actually administered. And access to ABEL's sophisticated causal models would be multiply advantageous. It would facilitate the extensions described above into etiological treatment, consistency-checking, and generation of explanations, as well as allowing a principled approach to unanticipated situations and enhancing the overall clarity of the program.

ABET has been an exploratory effort into the feasibility of using AI techniques in the domain of therapy for acid-base and electrolyte disorders. The performance of the program, while far from that of a human expert, is sufficiently good to establish the possibility of achieving genuine expertise, given the improvements outlined above. More importantly, we have devised a general approach to the formulation of acute symptomatic therapy, and demonstrated its applicability to at least one area. Whether our formalization will be as widely applicable as we hope remains to be seen.

References

- [1] Burke, Glenn, *LSB Manual*, MIT Laboratory for Computer Science, Technical Memo MIT/LCS/TM-200, 1981.
- [2] Fagan, Lawrence M., *VM: Representing Time-Dependent Relations in a Medical Setting*, PhD thesis, Stanford University, 1980.
- [3] Gorry, G. Anthony, Howard Silverman and Stephen G. Pauker, "Capturing Clinical Expertise: A Computer Program that Considers Clinical Responses to Digitalis," *The American Journal of Medicine*, 64:452-460, 1978.
- [4] Ingelfinger, Joseph A. and Peter Goldman, "The Serum Digitalis Concentration - Does it Diagnose Digitalis Toxicity?," *The New England Journal of Medicine*, 294:867-870, 1976.
- [5] Long, William, Shapur Naimi and M. G. Criscitiello, "A Knowledge Representation for Reasoning about the Management of Heart Failure," *Proceedings of the Ninth Annual Computers in Cardiology Conference*, 1982.
- [6] McDermott, John, "R1: A Rule-Based Configurer of Computer Systems," *Artificial Intelligence*, 19:39-88, 1982.
- [7] Moon, David A., *The Maclisp Reference Manual*, MIT Laboratory for Computer Science, 1978.
- [8] Patil, Ramesh S., *Causal Representation of Patient Illness for Electrolyte and Acid-Base Diagnosis*, MIT Laboratory for Computer Science, Technical Report MIT/LCS/TR-267, 1981.
- [9] Russ, Thomas A., *Ventricular Arrhythmia Management: A Knowledge-Based Approach*, MS thesis, Massachusetts Institute of Technology, 1983.
- [10] Sacerdoti, Earl D., *A Structure for Plans and Behavior*, Elsevier North-Holland, Inc., 1977.
- [11] Schwartz, William B., "Disorders of Fluid, Electrolyte, and Acid-Base Balance," in *Textbook of Medicine*, Beeson, McDermott and Wyngaarden, ed., W. B. Saunders Company, 1979.
- [12] Shortliffe, Edward H., *Computer Based Medical Consultations: MYCIN*, Elsevier North-Holland, Inc., 1976.
- [13] Szolovits, Peter, *Artificial Intelligence and Clinical Problem Solving*, MIT Laboratory for Computer Science, Technical Memo MIT/LCS/TM-140, 1979.
- [14] Szolovits, Peter, and William A. Martin, "Brand X: Lisp Support for Semantic Networks," *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, 940-946, 1981.
- [15] Weiss, Sholom M., Casimir A. Kulikowski, Saul Amarel and Aran Safir, "A Model-Based Method for Computer-Aided Medical Decision-Making," *Artificial Intelligence*, 11:145-172, 1978.
- [16] Winston, Patrick Henry, *Artificial Intelligence*, Addison-Wesley Publishing Company, 1977.