University of Pennsylvania
ScholarlyCommons

Technical Reports (CIS) | Department of Computer & Information Science

8-1-1990

TraumAID: Reasoning and Planning in the Initial Definitive Management of Multiple Injuries

Bonnie L. Webber
University of Pennsylvania, bonnie@inf.ed.ac.uk

John R. Clarke
Medical College of Pennsylvania

Michael Niv
University of Pennsylvania

Ron Rymon
University of Pennsylvania

María Milagros Ibáñez
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/cis_reports

Recommended Citation

Bonnie L. Webber, John R. Clarke, Michael Niv, Ron Rymon, and María Milagros Ibáñez, "TraumAID: Reasoning and Planning in the Initial Definitive Management of Multiple Injuries", . August 1990.

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-90-50.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/cis_reports/579
For more information, please contact repository@pobox.upenn.edu.

TraumAID: Reasoning and Planning in the Initial Definitive Management of Multiple Injuries

Abstract

The TraumAID system has been designed to provide computerized decision support to optimize the *initial definitive management* of acutely injured patients after *resuscitation and stabilization*. The currently deployed system, TraumAID 1.0, addresses penetrating injuries to the abdomen and to the chest. Our experience with TraumAID 1.0 has demonstrated some major deficiencies in rule-based reasoners that are faced with problems of both diagnosis and treatment. To address these deficiencies, we have redesigned the system (TraumAID 2.0), factoring it into two modules: (1) a *rule-based reasoner* embodying the knowledge and logical machinery needed to link clinical evidence to diagnostic and therapeutic *goals*, and (2) a *planner* embodying the global knowledge and logical machinery needed to create a plan that addresses *combinations* of goals. After describing TraumAID 2.0, we discuss an extension of the TraumAID interface (*critique mode interaction*) that may improve its acceptability in a clinical setting. We close with a brief discussion of management support in resource-limited environments, which is an important issue in the time-critical context of multiple trauma.

Comments

University of Pennsylvania Department of Computer and Information Science Technical Report No. MS-CIS-90-50.

This technical report is available at ScholarlyCommons: https://repository.upenn.edu/cis_reports/579

TraumAID: Reasoning and Planning In The Initial Definitive Management Of Multiple Injuries

MS-CIS-90-50
LINC LAB 180

Bonnie L. Webber
John R. Clarke
Michael Niv
Ron Rymon
María Milagros Ibáñez

Department of Computer and Information Science
School of Engineering and Applied Science
University of Pennsylvania
Philadelphia, PA 19104

August 1990

TraumAID: Reasoning and Planning in the Initial Definitive Management of Multiple Injuries

Bonnie L. Webber John R. Clarke Michael Niv Ron Rymon
María Milagros Ibáñez*

Running head title: TraumAID

Proofs to be sent to:
Bonnie Lynn Webber
Dept. of Computer and Information Science
University of Pennsylvania
Philadelphia PA 19104-6389

*Webber, Niv, Rymon and Ibáñez are members of the Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA 19104-6389, and Clarke is Director of the Trauma Center, Medical College of Pennsylvania, Philadelphia PA 19129. This research has been supported in part by ONR under subcontract grant N00129-89-C-0006 and by ARO under grants DAAG29-84-K0061 and DAAL03-89-C0031. The authors would like to thank Len Karpf, Charlie Ortiz and Mark Steedman for their valuable comments on earlier drafts of this paper. Questions and requests can be directed to bonnie@central.cis.upenn.edu.

Abstract

The TraumAID system has been designed to provide computerized decision support to optimize the *initial definitive management* of acutely injured patients *after resuscitation and stabilization.* The currently deployed system, TraumAID 1.0, addresses penetrating injuries to the abdomen and to the chest. Our experience with TraumAID 1.0 has demonstrated some major deficiencies in rule-based reasoners that are faced with problems of both diagnosis and treatment. To address these deficiencies, we have redesigned the system (TraumAID 2.0), factoring it into two modules: (1) a *rule-based reasoner* embodying the knowledge and logical machinery needed to link clinical evidence to diagnostic and therapeutic *goals*, and (2) a *planner* embodying the global knowledge and logical machinery needed to create a plan that addresses *combinations* of goals. After describing TraumAID 2.0, we discuss an extension of the TraumAID interface (*critique mode interaction*) that may improve its acceptability in a clinical setting. We close with a brief discussion of managment support in resource-limited environments, which is an important issue in the time-critical context of multiple trauma.

1 Introduction

Injuries, accidental and intentional, result in more years of human life lost in the United States than any other disease [1]. These injury-related deaths demonstrate a trimodal distribution [2]. The first peak of deaths occurs immediately as a result of lethal injuries and can only be eliminated by preventing the injuries themselves. The second peak of deaths occurs within the first hours of injury and can be eliminated or reduced by rapid delivery of expert care. The third peak is the result of late complications and can be ameliorated through expert care before and after the development of the complications. West [3] and others have clearly shown that, across the board, 30 to 40% of trauma deaths are preventable by the delivery of rapid and expert care.

Because of the need for rapid delivery of expert care, a major effort has been made to educate physicians so that they can provide an immediate expert response. This effort has been led by the American College of Surgeons through their Advanced Trauma Life Support (ATLS) course. The purpose of the course is to enable the physician to provide immediate expert response in the initial evaluation, resuscitation, and stabilization of severely injured patients [4]. The goal of our system, *TraumAID*, is to optimize the next step, their *initial definitive management.*

Such an effort is warranted: the ATLS course assumes that definitive expert care will be available through communication or transportation after the patient has been resuscitated and stabilized during the initial hour of care. But that is not always the case. Many parts of the United States have no trauma care delivery systems and many hospitals are so remote from trauma centers that transportation to trauma centers is not always feasible. And the situation is not improving: the *New York Times* recently reported a large drop in the number of hospitals offering trauma care:

> Around the country, hospitals that offer trauma care for victims of life-threatening stabbings, falls, gun shots or accidents are closing to ambulances. Los Angeles County has seen 12 of the 20 hospitals offering trauma care do so in two years. St. Louis, Miami and Detroit have made similar cuts, and New York has severe overcrowding.
>
> *New York Times*, 22 October 1989, p.22

Even where trauma care facilities are available, the quality of care can vary. For example, since trauma presents primarily during nights and weekends, experts may be off-duty or compromised by fatigue. The result is that expert support cannot be guaranteed. By helping to prevent errors in the initial steps of definitive management, TraumAID aims to prevent complications, thereby reducing the second and third mortality peaks for injured patients, which account for 50-60% of deaths in trauma cases.

Most decision support aids for trauma have to date been simple scoring algorithms [4,5]. The most widely used aid is the *Trauma Score* developed by Champion and Sacco [6], which itself incorporates the Glasgow Coma Score, a scoring system for head injuries [7]. These simple systems are not diagnostic tools: rather, they correlate with prognosis and are used for triage decisions. The complexity of the decisions involved in managing patients with complex multiple injuries requires more than simple scoring systems can provide. TraumAID attempts to provide such support using Artificial Intelligence techniques of reasoning and planning. TraumAID 1.0 currently provides decision support in the initial definitive management of patients with penetrating injuries of the abdomen and chest. It has already undergone preliminary validation [8,9], and further validation studies are being conducted at the Medical College of Pennsylvania Trauma Center.

In this paper, we first provide an overview of the original rule-based system TraumAID 1.0 and its mode of operation (Section 2). An example illustrating a physician's interaction with this version of the system appears in Section 3. The core of the system is described in more detail in Section 4, with comments on its current level of performance given in Section 5. Our experience with TraumAID 1.0 has demonstrated some major deficiencies in rule-based reasoners faced with problems of both diagnosis and treatment. To address these deficiencies, we have redesigned the system, factoring it into two modules: (1) a *rule-based reasoner* embodying the knowledge and logical machinery needed to link clinical evidence to *goals*, and (2) a *planner* embodying the global knowledge and logical machinery needed to address *combinations* of diagnostic and therapeutic goals, recommending to the physician how best to address them. This new system (TraumAID 2.0) is described in Section 6. Section 7 describes our work on providing an alternative mode of interacting with the system (*Critique Mode Interaction*), which may improve its acceptability in clinical settings. Finally, we close with a brief discussion of providing management support in resource-limited situations.

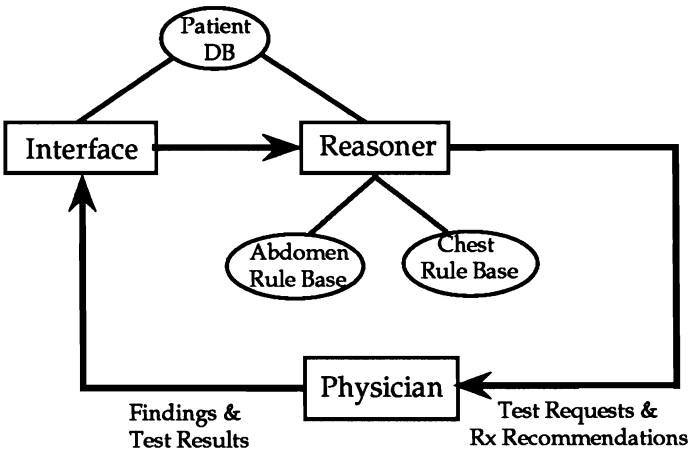

Figure 1: TraumAID 1.0 System Diagram

2 Overview

TraumAID is the current result of an on-going multi–year collaboration between the director of the regional trauma center at the Medical College of Pennsylvania and members of the Department of Computer and Information Science at the University of Pennsylvania. Its long-term goal is to provide computerized decision support to optimize the *initial definitive management* of acutely injured patients *after resuscitation and stabilization.* Currently, TraumAID addresses penetrating injuries to the abdomen and to the chest. Extensions are planned for penetrating injuries to the upper extremities, lower extremities, head, neck and perineum, and then for blunt injuries to the same areas.

TraumAID is intended for otherwise healthy, adult patients. While it does not currently deal with pregnant women who have suffered abdominal or chest injuries, or such injuries to patients with other significant medical conditions, it does accommodate patients with problems of drug and/or alcohol intoxication, which are often present in trauma patients. This covers a large portion of trauma patients: at the Medical College of Pennsylvania, approximately *one quarter* of all patients hospitalized on the trauma service have penetrating injuries, 99% of which are gunshot wounds or stab wounds, 90% of which occur in patients under the age of 55, the age which the American College of Surgeons has identified as the point at which age begins to influence outcome. 95% of those under the age of 55 (that is, 85% of patients with penetrating injuries) have no medical problems of signficance other than intoxication.

The core of the basic system (TraumAID 1.0) is a rule-base expert system, initially written in LISP to run on a Symbolics workstation, and subsequently implemented in C as well, to run

on a PC. A diagram of the basic system is shown in Figure 1. Its basic mode of operation is as follows:

Using windows and menus, TraumAID 1.0 accepts an initial description of the wound(s) in terms of their type, location and direction, as well as any initial findings (positive or negative) that the physician may report. It then reasons *forward* from findings to *diagnostic suspicions* warranting further diagnostic investigation and to *diagnostic conclusions.* This phase of forward reasoning is followed by a phase of *backward reasoning* to determine what evidence is needed in order to confirm or rule out its current suspicions. Whenever a treatable diagnosis is concluded, the system accesses its set of suggested therapeutic procedures, which are eventually packaged together into the system's suggested management plan.

If the system lacks the information needed to confirm or rule out a suspicion, it requests this information from the physician – symptoms and findings first, then information provided by diagnostic tests. The physician can choose which requests to address. Any information provided will trigger further forward chaining to suspicions and conclusions, so that this basic cycle may be repeated several times. Note that since it is possible for medical procedures to be both therapeutic and diagnostic, the outcome of treatment itself may trigger a new round of suspicions and conclusions. In this basic cycle, TraumAID 1.0 can consider 101 findings to reach 61 conclusions, through 454 decision rules. *Certainty factors* [10] are not used in this process, as (according to co-author Clarke) experts in the acute care of multiple injuries, on whose reasoning the system's is based, tend to reason categorically, using protocol-like rules, to avoid time-consuming contemplation.

The system has a part-whole hierarchy representing the anatomy of the abdomen and of the chest, to allow rules to be stated more succinctly and to reduce their number. The system can distinguish right and left, so that something can be true of the right lung and false of the left. The system can handle iterations – for example, to accommodate the fact that a pneumothorax can recur or persist after treatment by a chest tube. Finally, it has a primitive facility for recognizing and accommodating interactions between two recommendations, using fixed diagnostic sequences and additional *ad hoc* rules. TraumAID 2.0 replaces this primitive coordination facility with an incremental planner, as discussed in Section 6.

There are currently two versions of TraumAID: one designed for use by physicians in a well-equipped Emergency Center, the other for use by independent-duty medical corpsmen on submarines. (The latter was developed under a contract with the Naval Submarine Medical Research Laboratory.) The two versions are not independent: the submarine version was developed as a subsequent modification of the Emergency Center system. Features of the submarine version of TraumAID and its relationship to the original version are discussed in Section 8.

3 Example

The following example (based on one of Clarke's actual cases) illustrates the feel of a very simple interaction with TraumAID 1.0.

> Patient WR was a healthy male, who was shot from the left side with the bullet entering the left lower quadrant of the abdomen. There was no exit site. The patient had abdominal guarding, but no tenderness or rebound tenderness. There were no other abnormal findings on physical exam.

Figure 2 shows a snapshot of the system screen after these initial findings have been entered, and the physician has turned control of the interaction over to the system (a mode of interaction called *Let System Ask*). At this point, the system has asked about two additional findings, shock and unconsciousness, and is just asking about a third, obtundation.

Figure 3 shows the system continuing its questions, asking about evisceration, peritoneal scarring, and weak pulse. Questions involving diagnostic tests are not asked until after those that simply require examination. Here the system asks about the results of urinalysis (checking for blood in the patient's urine) and of a cystogram. Notice that as well as "yes" and "no", the physician may indicate that he or she does not currently know the information by answering "unknown". At this point, there is no more information that the system takes to be necessary for confirming or ruling out its current suspicions. Forward reasoning leads it to the conclusions shown in Figure 3 and the treatments recommended for them. (Treatable conclusions are shown in boldface.)

Figure 4 demonstrates the system's simple explanation facility: the physician can ask for justification of any diagnostic conclusion or treatment recommendation (through the command *Why*), and a response will appear on the right-side of the screen. Justification of treatments is given in terms of the diagnoses that motivate them. Justification of diagnoses are given in terms of the rules by which they have been concluded. (More detailed justification of treatment recommendations are given when the system is in *Critique Mode*, cf. Section 7.)

4 System Description

4.1 Rules

Here we describe TraumAID 1.0's rule-based reasoning in more detail – first, its rules and then its cycles of reasoning. TraumAID 1.0 uses two types of rules in its forward and backward reasoning: *suspect rules* and *conclude rules*. Suspect rules are used solely in forward reasoning, to

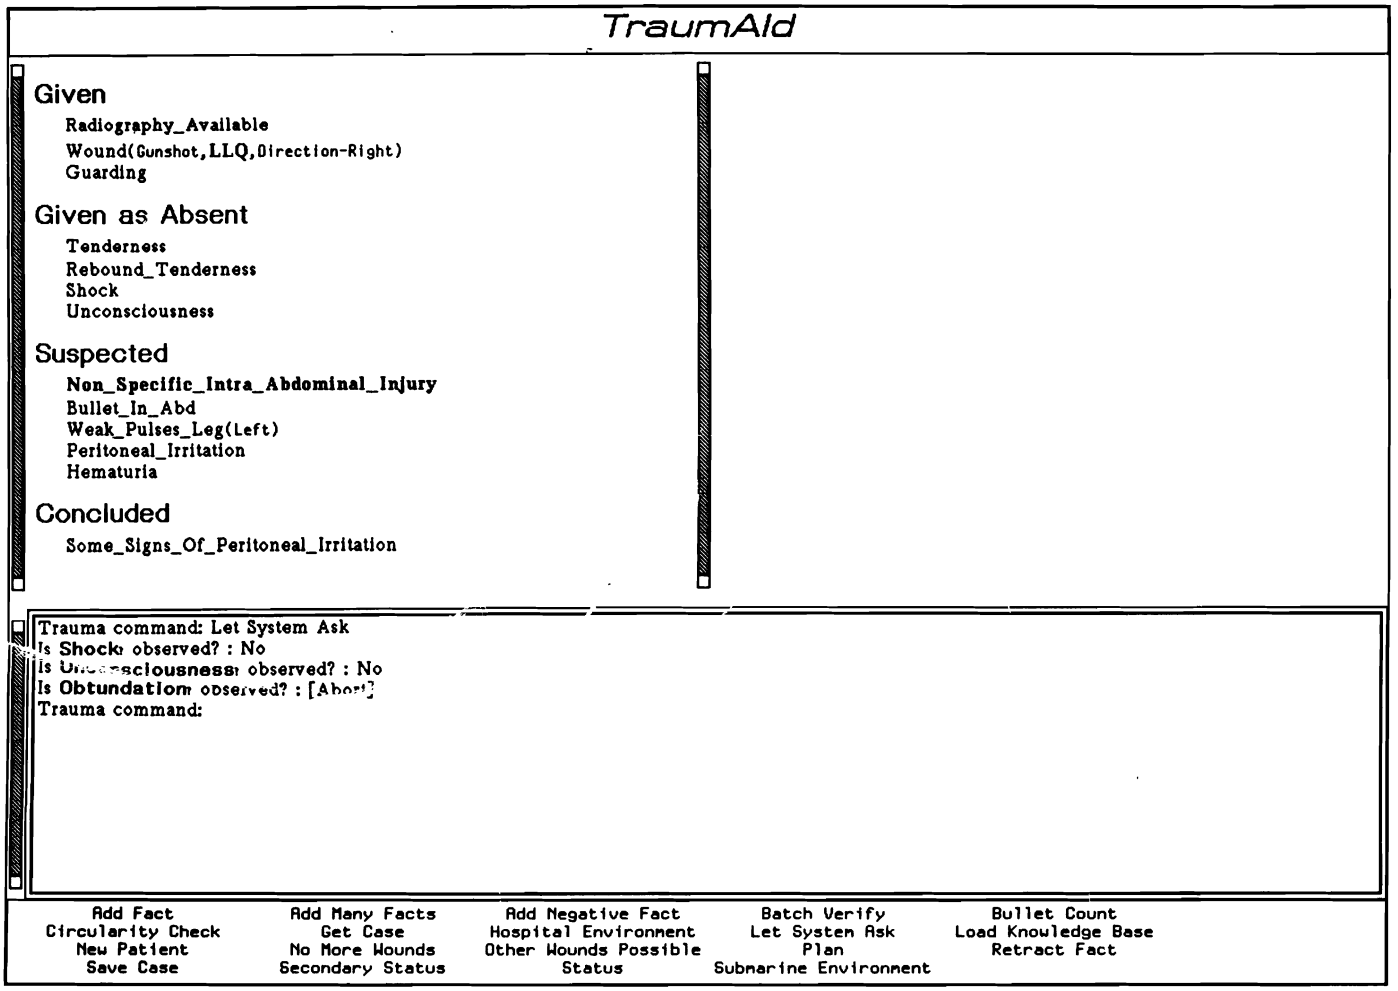

Figure 2: TraumAID 1.0 Initial Screen

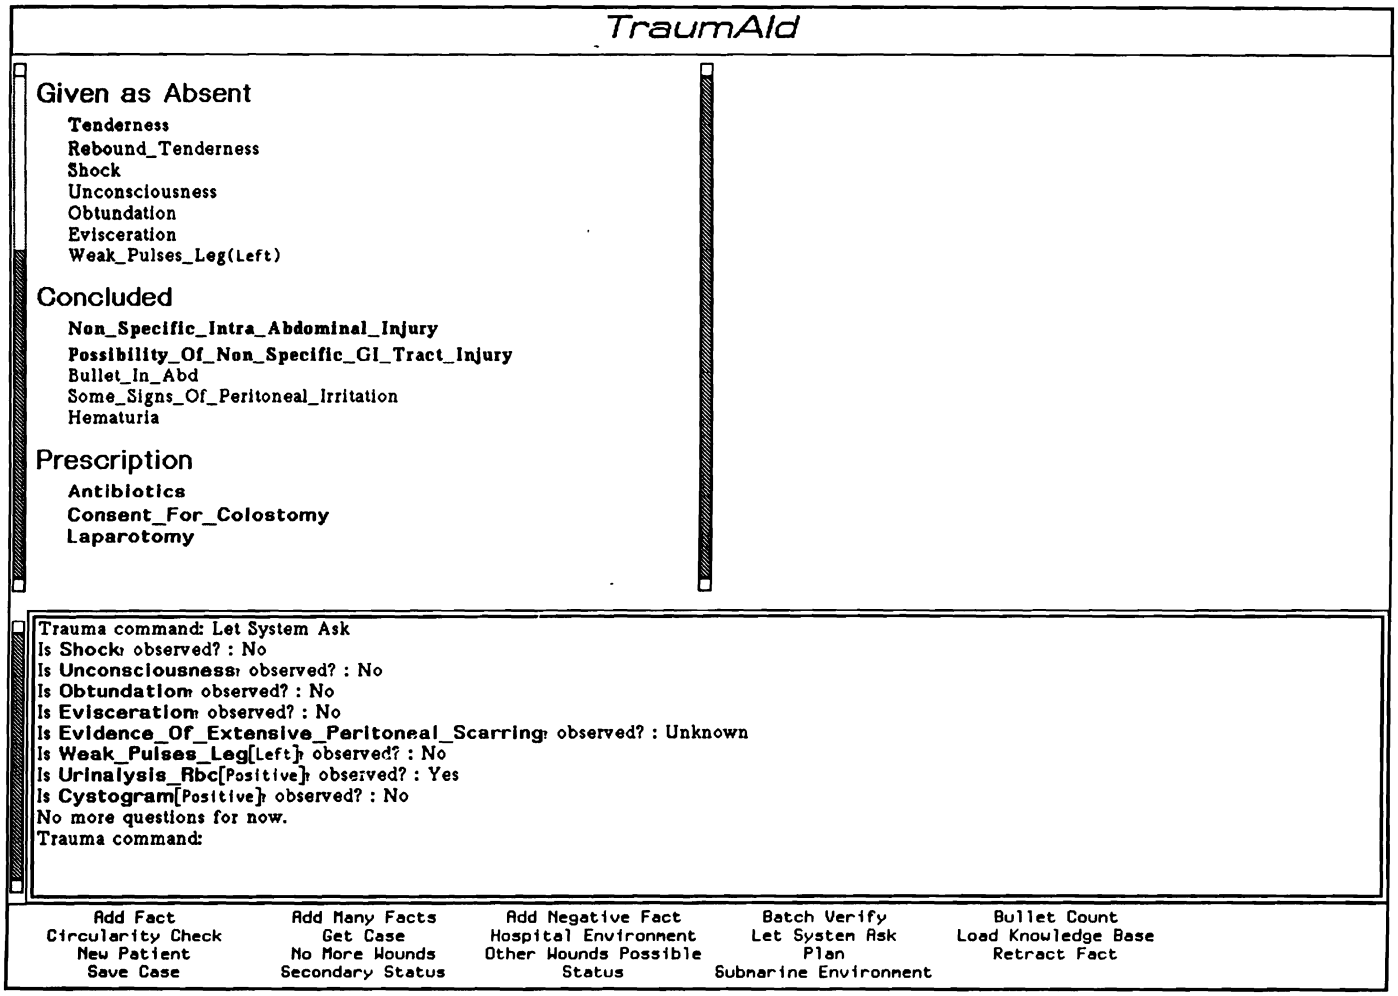

Figure 3: TraumAID 1.0 Final Screen

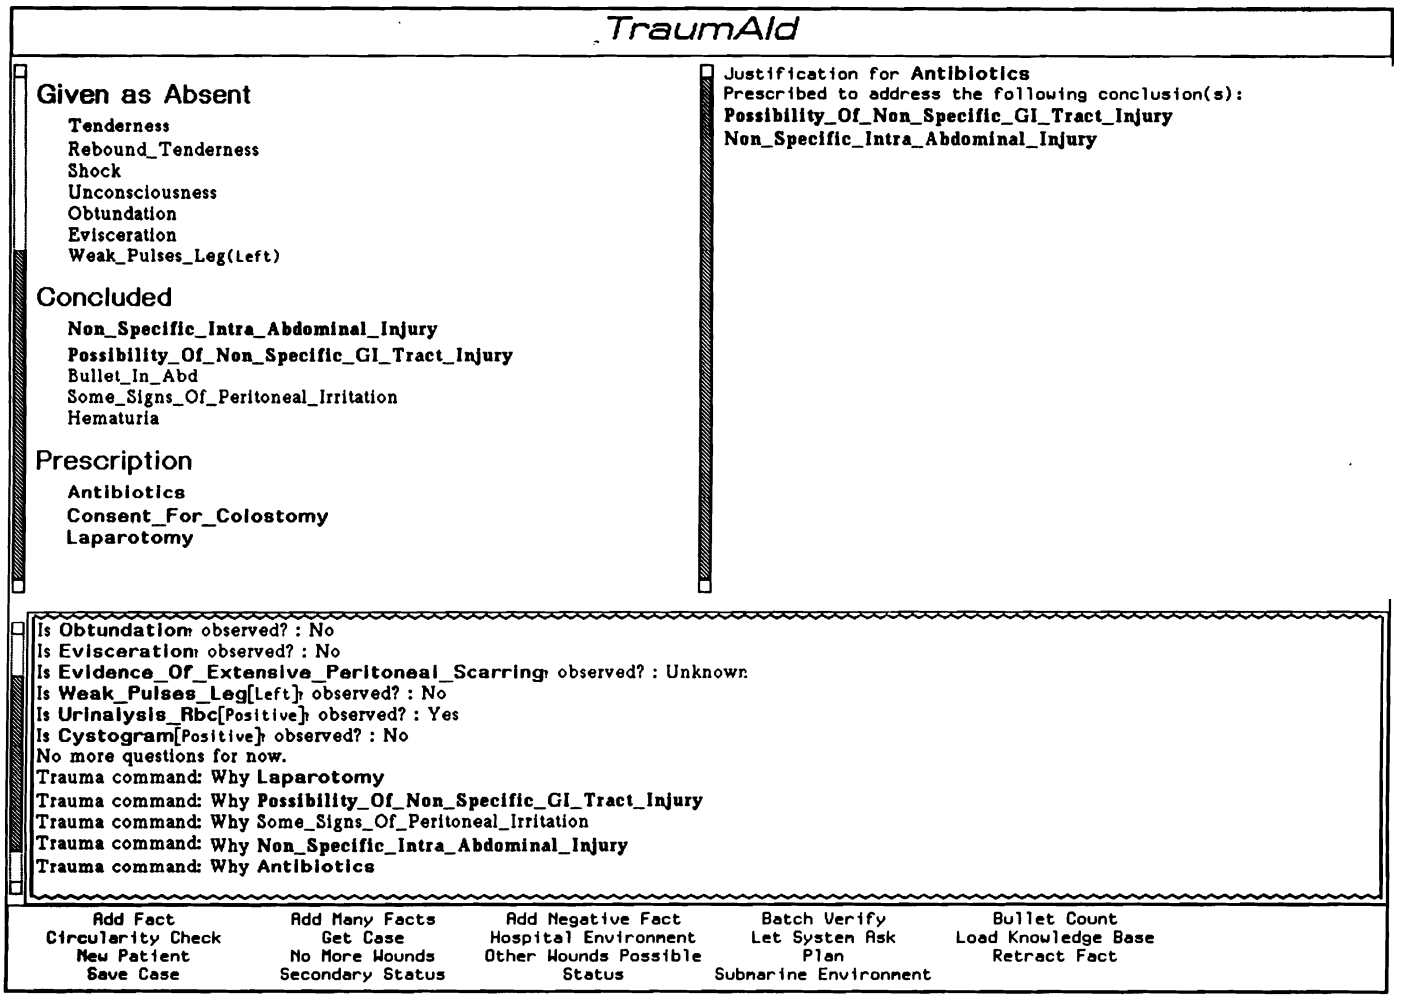

Figure 4: TraumAID 1.0 Explanation Facility

link findings, test results and the fact that particular actions have been performed to *suspicions* about the current state of the patient. For example, the following rule leads the system to suspect that the patient may be suffering from a pericardial tamponade if he has a chest wound and his heart sounds are muffled. (Rules use a Prolog-like syntax. *Suspect rules* are indicated by the "?" after the rule number.)

5405? pericardial_tamponade :-
    wound(location = 'chest),
    muffled_heart_sounds.

Once the system suspects a particular diagnosis, it uses its corresponding *conclude rules* to check on evidence during the next phase of backward-chaining. **Rule 5412** illustrates one rule for concluding pericardial tamponade.

5412 pericardial_tamponade :-
    muffled_heart_sounds,
    ultrasound_effusion(result = 'positive),
    -shock.

If pericardial tamponade is suspected, and the patient has muffled heart sounds and is not known to be in shock, then if the result of the diagnostic test *ultrasound_effusion* is not currently known, the system is led to request the test. This is how TraumAID requests that the physician take particular *diagnostic actions.*

The reason for emphasizing *known* in specifying truth values is that TraumAID rules are written using a four-valued logic, in which one can require atomic clauses to be known true or known false (-), or *not known* to be false (%), or *not known* to be true (-%). For example, the third clause of **Rule 5412** above specifies that the patient is *known not* to be in shock. Entries in the Patient Database (cf. Figure 1) are *three-valued*, indicating that a particular finding or test result is either positive, negative or unknown. Thus the *rule-language* values of *not known true* and *not known false* correspond to disjunctions of database values. For example, the third clause of **Rule 5400** below specifies that the patient is *not known to have* a distended abdomen. It would be supported by either "distended_abdomen=unknown" or "distended_abdomen=negative" in the patient DB.

5400? pericardial_tamponade :-
    wound(location = 'epigastric, not direction = 'down),
    shock,
    -% distended_abdomen.

The system suspects pericardial tamponade if the patient has a wound of the epigastrium and is in shock, unless it is known that the abdomen is distended (in which case, an abdominal injury is a more likely explanation for the shock than pericardial tamponade).

Having been used during backward reasoning to identify the evidence needed for drawing conclusions, conclude rules are subsequently used in forward reasoning to actually draw those conclusions. They are also used to generalize conclusions and thus enable rules to be stated more succinctly. Currently, this is how anatomical and abstraction relations among injuries are specified. For example, the following conclude rules are among those used in forward reasoning to generalize specific chest injuries to the concept *penetrating_chest_injury* (along with the side that the injury is on).

5127 penetrating_chest_injury(side = S) :-
simple_pneumothorax(side = S).

5128 penetrating_chest_injury(side = S) :-
tension_pneumothorax(side = S).

5129 penetrating_chest_injury(side = 'left) :-
pericardial_tamponade.

This in turn enables the system to specify a single rule like **5108**, instead of a set of rules, one for each specific injury.

5108 lacerated_diaphragm(side = S) :-
wound(location = 'abdom),
penetrating_chest_injury(side = S),
-wound(location = 'chest).

The system concludes that the patient's diaphragm has been injured on either the left or right side if the patient has sustained an abdominal wound, is suffering some type of penetrating chest injury on that side, but doesn't have a wound in his chest.

Forward reasoning also employs a form of "closed world" negation. That is, if one of the antecedents of a conclude rule is falsified, the rule is too. If all the conclude rules for a proposition P are falsified, then P itself is concluded false. The "closed world" assumption here is that there is no other way of concluding the truth of P. (Note that our use of a four-valued logic means that this is not the same as Prolog's "negation as failure" reasoning. A proposition cannot be falsified simply on the basis of unknown information.)

A simple example of TraumAID's "closed world" negation involves the condition *moribund*:

5335 moribund_patient :-
shock,
unconscious.

Because this is the only conclude rule for moribund in our domain, the system will conclude

that a patient is moribund if he is in shock and unconscious. If either condition is known not to hold, then it will be concluded that the patient is not moribund. Finally, if only one condition is known to hold and nothing known of the other, then no conclusion will be drawn.

Finally, whenever a treatable diagnosis is concluded by virtue of all the clauses in an appropriate *conclude rule* being true, TraumAID reasons on its next phase of forward-reasoning to appropriate therapeutic actions. This is how TraumAID recommends that the physician take appropriate *therapeutic actions*, for example:

treat pericardial_tamponade with

continuous_decompression,
operative_exposure_of_heart,
heart_repair.

As in clinical medicine, TraumAID was designed to do no more diagnosis than is called for by the needs of treatment. Its rules were constructed by identifying first the available treatments, then the diagnostic conditions, and only then, the findings and test results that would lead one to make these diagnoses. Thus TraumAID is not subject to the criticisms of diagnostic systems presented in [11].

4.2 Cycles of Operation

The fact that reasoning (phases of forward and backward chaining) by the system alternates with performance by the physician gives the system a great deal of power, which it uses in three ways: (1) to embed an explicit process of evidence gathering; (2) to enable it to consider the outcome of actions, be they diagnostic or therapeutic or both; and (3) to enable reasoning about complications of treatable diagnoses.

4.2.1 Gathering Evidence

We mentioned in Section 2 that TraumAID 1.0 does not use *Certainty Factors* in deciding whether it has sufficient evidence to justify a particular diagnostic conclusion. Instead, it makes use of *protocols* to focus its pursuit of suspicions. That is, a clue to a diagnosis will prompt collection of other clues, and when the easily obtainable clues justify it, a definitive (and usually more invasive or costly) test will be recommended. These protocols are embedded in the system's cycles of operations. For example, suppose during a phase of *forward reasoning*, the system concludes on the basis of any of several wound types (including a wound to the left upper parasternal region that is not directed left) or a finding of *hemoptysis*, that the patient has (what we label)

Slight_Possibility_of_Tracheal_Injury

This is one of several pieces of evidence, including also a potential mediastinal injury, that will lead the system to suspect (during its following phase of *forward reasoning*) that the patient has (what we label)

Possibility_of_Tracheal_Injury

It will then find during *backward reasoning* that it can confirm this suspicion through several types of evidence. It thus asks the physician whether the patient demonstrates *stridor* or whether pneumomediastinum is found on X-Ray. If the physician reports positive evidence of either of these, the system will then be led to suspect

Tracheal_Injury

(It can also be led to suspect tracheal injury directly, on the basis of a bullet in the patient's superior or posterior mediastinum, or joint findings of *hemoptysis* and *stridor*.) To enable it to finally conclude or rule out tracheal injury, the system will request a bronchoscopy be done. If the physician then reports positive test results, tracheal injury is concluded and therapeutic procedures are recommended. This regular step-by-step pursuit of evidence to either confirm or rule out tracheal injury is one of the protocols embedded in cycles of system reasoning followed by actions on the part of the physician.

4.2.2 Responding to Outcomes

One of the interesting features of certain actions taken in managing multiple trauma is that they may be either therapeutic or diagnostic, depending on their outcome. For example, a *tube thoracostomy* can be used to *treat* a simple hemothorax if it is successful in draining accumulated blood from the thoracic cavity. If it does not succeed (i.e., blood drainage is either massive or persistent or both), it is *diagnostic evidence* that the patient is suffering from a massive hemothorax instead, which requires different treatment. Only by returning to a phase of reasoning following the performance of actions can the system check on the results of recommended actions and interpret their diagnostic yield.

The system also uses its cycles of reasoning and performance to avoid both the assumption that actions will be done correctly and the assumption that equipment will function properly. For example, after recommending a *primary tube thoracostomy* (for either a simple hemothorax or a simple pneumothorax), when the system is next invoked, it will be led to suspect that the tube has been misplaced or that it is malfunctioning. To rule out (or confirm) these suspicions,

the system requests that a follow-up X-ray be done. If either condition is then reported to apply, the system recommends it be treated with a *replacement tube thoracostomy*. (The system does not assume the replacement tube will be inserted correctly, nor that it will function properly. It keeps recommending replacement tubes until the procedure is done correctly and the tube is found to be working properly.) Again, it is only by returning to a phase of reasoning following the performance of actions that the system can check on whether actions have been done correctly and their results worth interpreting.

4.2.3 Dealing with Complications

In trauma, as other areas of medicine, injuries can cause complications. The original injury must be treated, while suspicions of further injuries must be pursued. For example, if the results of an arteriogram lead the system to conclude that the patient has an injury to his descending aorta, the system will recommend that a *thoracotomy* be done, in order to repair the injury. However, it will also be led to suspect that the patient may have sustained injury to his spinal cord (*ischemic spinal cord injury*) as a result of reduced blood supply. The system seeks evidence of motor loss in both of the patient's legs in order to confirm or rule out this suspicion, which may involve a new request to the physician for information.

While complications *can* be handled in a simple expert system (that is, one in which a single phase of diagnostic reasoning is followed by a single phase of therapeutic reasoning), such a system cannot make any therapeutic recommendations known until all its diagnostic conclusions have been reached. A system that cycles is not so limited.

4.3 Other Features

As noted earlier, after the system identifies to the physician the diagnostic and therapeutic procedures it recommends, it will recommence reasoning based on whatever information the physician provides. That is, it does not *insist* that its recommendations be carried out. However, because of the way that its rules are written – that is, to facilitate its identifying (during backward reasoning) the procedures it should recommend – TraumAID 1.0 is limited in its ability to make use of all the *unsolicited* information that a physician might provide. The following example illustrates the problem: suppose a patient presents with a chest wound and is found to be in shock. The system will be led to suspect a *pericardial tamponade*. Normally the best diagnostic test for this condition in a patient in shock is *needle aspiration of the pericardial sac*, since this procedure can reduce pressure on the heart and, for a short time, the shock as well. If the patient is not in shock, *ultrasound effusion* is used because it is less invasive. TraumAID's

rules are written (see **5410** and **5411** below) in a way that directs its recommendation to needle aspiration for patients in shock and to ultrasound effusion otherwise. Thus if a physician happens to have done an ultrasound effusion, for other reasons, on a patient in shock, TraumAID 1.0 is blocked by the *-shock* antecedent in Rule **5411** from making use of the results. By Rule **5410**, it can only make use of that test if the patient is not in shock.

5410 pericardial_tamponade :-
...,
shock,
needle_aspiration_pericardial_sac(RESULT = 'positive),
...
....

5411 pericardial_tamponade :-
...,
distended_neck_veins,
ultrasound_effusion(RESULT = 'positive),
-shock,
...
....

This limitation has been addressed in the extended system TraumAID 2.0 (cf. Section 6.1.2).

5 Performance of TraumAID 1.0

In developing TraumAID 1.0, we have carried out both continuous verification during rule-base development and prototype *post hoc* validation tests after the incorporation of each module. In addition, special verification was done of the Submarine version of the system to make sure that its recommendations were in line with standard Navy practices.

Our continuous verification procedure makes use of a data base of 399 theoretical cases, on which the system checks itself in *batch validation* mode. Whenever its conclusions or recommendations for a case change as a result of changes in its decision rules, complete information is printed out on those cases, for further analysis.

The first module developed for TraumAID was in aid of penetrating injuries of the abdomen. It could consider 56 findings to reach 16 conclusions, using 122 decision rules. A prototype validation test was performed on this module [8] in which TraumAID's recommendations were compared with those of physicians-in-training for five representative patient presentations at the Medical College of Pennsylvania trauma center. These trainees ranged from third-year surgical clerks to chief surgical residents. TraumAID's recommendations, those of each of the trainees, and the actual care were abstracted, blinded, and then judged by surgical faculty not otherwise involved in the project. Only the actual care and the recommendations of TraumAID were judged to be acceptable for all five cases. The average rank of TraumAID's recommendations

was higher than that of any of the chief surgical residents (although not significantly better).

The second module developed for TraumAID deals with penetrating injuries to the chest. A similar prototype validation test was performed on the extended system. Here, compared with the recommendations of three more recent chief residents at the trauma center on ten representative, but diverse cases involving penetrating chest injury, TraumAID's recommendations were judged better than twenty of the residents' plans and worse than only seven [9] ($p < 0.05$). Moreover, while all of TraumAID's recommendations were judged acceptable (one was judged deficient but acceptable), four resident plans were judged unacceptable.

In November 1989, a PC-based version of TraumAID 1.0 was installed in the Trauma Resuscitation Unit at MCP. Since then the system has been consulted in all cases of penetrating wounds to the chest and/or abdomen. A preliminary review of the first 50 cases indicates that the system can be used for most cases as intended and the the advice is usually acceptable and some time helpful.

This PC-based version of TraumAID 1.0 will remain in the Trauma Resuscitation Unit at MCP, both to expand its database of actual cases and to allow us to determine its strengths and weaknesses in actual clinical use. To see what can be gained through "user-friendliness", we have designed a HyperCard interface for TraumAID that will be deployed in parallel this fall. Users' responses to both interfaces will be compared for ease of use and correctness of data entries.

6 Planning: Between Reasoning and Acting

6.1 General Issues

6.1.1 Procedure Interactions

One of the major problems we have had to address in developing TraumAID is the fact that it is not uncommon for patients to present with multiple problems: some may be caused by the same injury, others may result from distinct injuries. Each problem demands its own diagnostic and therapeutic procedures, which the physician in charge must follow. Taken together, they present a major problem of what the physician should do and when he or she should do it. Moreover, because of potential interactions between these procedures, it is not simply a sequencing problem: what is an appropriate procedure for an isolated injury may not be appropriate in the context of other problems that must be addressed. (Similarly, a procedure that is less preferred for an isolated injury may be more preferable in the context of other injuries.)

TraumAID 1.0 mediates the most common of these procedure interactions in a distinctly non-modular fashion, through a set of *ad hoc* rules. For example, when the system concludes that the patient is suffering a left carotid artery injury it recommends

treat l_carotid_artery_injury with

operative_exposure_mid_superior_mediastinum,
l_carotid_arterial_injury_repair.

This is actually not a complete specification of what needs to be done, because the particular *method* of operative exposure depends on what other actions are being recommended. What happens is that the system generalizes (cf. Section 4.2) a wide variety of injuries (including upper and lower thoracic esophageal injuries, tracheal injury, bronchial injury, massive or persistent hemothorax, and injury to the descending thoracic aorta) to a single injury type labelled **condition_requiring_thoracotomy**. This results in two separate rules to select an operation to treat left-carotid injuries:

5969 exposure_thru_L_thoracotomy :-

l_carotid_artery_injury,
condition_requiring_thoracotomy.

5988 exposure_thru_median_sternotomy :-

l_carotid_artery_injury,
-% condition_requiring_thoracotomy.

That is, if the patient has a left carotid artery injury and is ***known*** to need a thoracotomy (either side), then the system recommends that the artery injury be treated with a left thoracotomy. (Of course, the operation must address both needs, not just the artery injury.) On the other hand, if the patient is ***not known*** to need a thoracotomy, then the system recommends treating the

artery injury with a median sternotomy. While on a small scale, one may be able to anticipate all such interactions and augment TraumAID's rules in order to deal with them, this solution is impractical as system coverage expands to other parts of the body and to blunt injury as well.

The solution we have chosen to the problem of developing an orderly and consistent management plan is to factor TraumAID into two linked components:

- a *rule-based reasoner* embodying the knowledge and logical machinery needed to link clinical evidence to *goals*: (1) *diagnostic goals*, to confirm or allay diagnostic suspicions through knowledge acquired in the performance of diagnostic procedures, and (2) *therapeutic goals*, to treat concluded diagnoses through the performance of therapeutic procedures;
- a *planner* embodying the global knowledge and logical machinery needed to address *combinations* of diagnostic and therapeutic goals, recommending to the physician which to address at the current stage of patient management and how best to address them.

Many of the characteristics of this planner follow from the same domain features that have led to the cycles of reasoning and acting used in TraumAID 1.0. In particular, since diagnostic procedures are as much a part of plans as therapeutic procedures, the planner cannot assume that it knows its final therapeutic goals from the very start. (The point of diagnostic procedures is clearly to identify, directly or indirectly, therapeutic goals that the plan must then address.) Planning in this context requires that the results of plan execution (e.g., new information, changes in the patient's state, etc.) feed a new round of reasoning and goal setting, until a final therapeutic goal is recognized and appropriate therapeutic procedures recommended. In the initial definitive management of multiple trauma, this happens when the patient is sent to the operating room or trauma unit or is discharged.

6.1.2 Physician Accommodation

We noted in Section 4.3 that one failing of TraumAID 1.0 was its limited ability to make use of unsolicited test results. In TraumAID 2.0, the separate roles of rule-based reasoner and planner ensure that the reasoner can make use of any evidence, independent of what led to its acquisition.

6.1.3 Plan Quality and Plan Execution

There are two other general points we want to discuss before briefly describing the operation of our planning component. These relate to (1) the *quality* of the plan created by the planner and (2) *how much* of it is executed before the next phase of reasoning commences. The two points are not independent: one's decision vis-a-vis the former is clearly contingent on one's decision vis-a-vis the latter, which we therefore address first.

At one extreme, one may assume that *all* the actions recommended in the plan are executed before reasoning recommences. However, at least two things can happen in the domain of multiple trauma that make this option unattractive:

1. A change in the patient's state or a diagnostic action may reveal a condition requiring immediate major surgery, thereby terminating *Initial Definitive Management*. If *all* recommended actions were performed before the next phase of reasoning commences, costly and perhaps risky actions may have been done unnecessarily. (N.B. Not all diagnostic or therapeutic goals need be satisfied by the end of *Initial Definitive Management*: some may have to be delayed until the third (post-operative) phase of trauma care.)

2. Diagnostic suspicions are not necessarily independent: one suspicion, confirmed or ruled out, may effectively rule out another. Thus it may not be *necessary* to perform all independently suggested diagnostic tests. For example, given a knife wound to the chest below the level of the diaphragm (and no other wounds), the system will be led to suspect a *lacerated diaphragm* (**Rule 5101**). If the patient is in shock, the system will also be led to suspect both a *non-specific abdominal injury* (**Rule 6315**), as well as a *massive hemothorax* (**Rule 5280**). Note that these suspicions are not independent: for a knife wound to have caused abdominal injury, it must penetrate the diaphragm. Moreover, suspicions of lacerated diaphragm and massive hemothorax are also not independent, since one requires the knife to have been pointing up, the other, down.

In both cases, performing *all* the procedures recommended in a plan can lead to unnecessary risk to the patient, as well as increased medical costs.

The other extreme is carrying out only the *first* procedure in the plan, before returning to consider the new information resulting from the procedure and the patient's new state. The problem with this is that it can end up wasting the effort that went into forming a plan in the first place. This may have involved a great deal of effort, which would be the case if the plan had been optimized to satisfy all the goals posed by the reasoning component.

This brings up the point about *plan quality*. Optimizing a plan – that is, minimizing its cost in resources, time and/or money, *over and beyond* avoiding actions or sequences of actions that put the patient at additional risk – is computationally expensive. If one is not going to execute an entire plan, globally optimizing it appears unnecessarily wasteful, given the time-critical nature of the enterprise.

The above is not an unfamiliar situation. It is at the core of much of the work in *reactive* or real-time planning: one must decide both how much effort to invest in forming a plan and, subsequently, how much to recommend carrying out before reflecting on the results. Where our

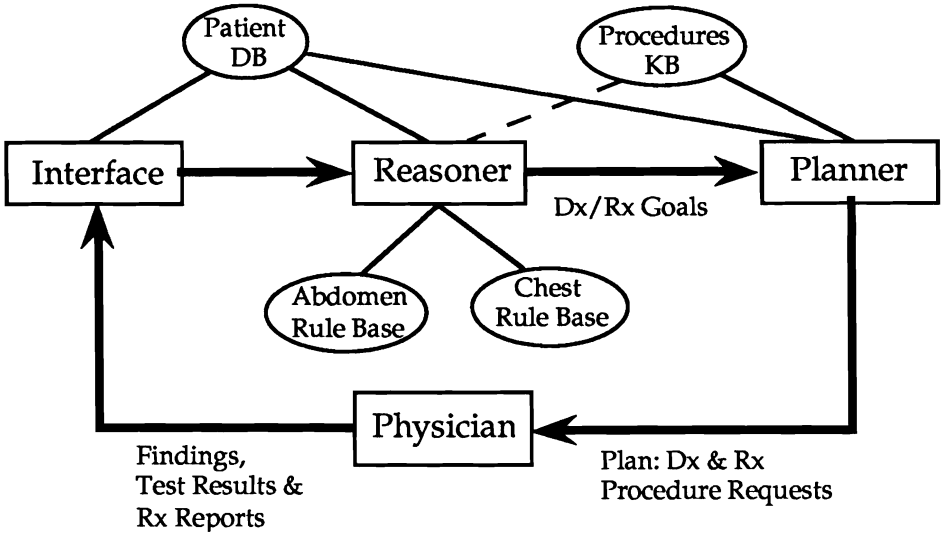

Figure 5: TraumAID 2.0 System Diagram

work differs from work in reactive planning is that many of the agent's (i.e., the physician's) planned actions are meant to comprehend the situation, not to directly rectify it.

The planner we have designed for TraumAID 2.0 walks a middle ground: it develops a plan that tries to satisfy all its given goals in an order that adheres to *a priori* domain constraints, but only invests additional resources to optimize the initial segment. Thus, even though the physician may only carry out one or two procedures before reporting their results back to the system, he or she is always given a global view of what the patient's *current* situation may require.

6.2 The TraumAID 2.0 Planner

Here we describe the planner being incorporated into our revised system, TraumAID 2.0. Since the planner is described elsewhere in more detail [12,13], the description here will be brief. A diagram of the augmented system is shown in Figure 5.

In TraumAID 2.0, the rule-based reasoner presents the planner with a set of goals (goals of acquiring information or of treating injuries). As noted, these may change over the course of plan execution. In response, the planner must develop a partially ordered set of procedures that satisfy these goals. It is expected that the physician will do at most the first one or two of these before reporting results back to the system, allowing it to recommence reasoning.

What we would like to illustrate is: (1) the knowledge that TraumAID 2.0 uses in mapping goals to procedures, and (2) the method by which TraumAID 2.0 comes up with a plan that

satisfies all the goals posed by the reasoner, or if that is not possible because of conflicts or resource limitations, a plan that satisfies the most urgent and important goals.

6.2.1 Procedure Knowledge Base

TraumAID 2.0's *Procedure Knowledge Base* (PKB) embodies knowledge of *management goals*, *management procedures* and the relationships between them. Significantly, the relationship is many-to-many, in that a goal may be satisfied by more than one procedure and a procedure may be used to satisfy more than one goal.

With respect to goals, the PKB contains information about their *type*, *urgency*, *related procedures* and *priority*. The type of a goal is either diagnostic (**RO**) or therapeutic (**Rx**). A goal is considered urgent if it addresses existing shock or a cause of existing shock (e.g. massive blood loss). A goal can be satisfied by members of its list of related procedures. Their order in the list reflects their preference ordering in the case of isolated injury. With respect to procedures, the PKB contains information about their their *performance sites*, their *preconditions* and the *goals* they satisfy. A particular instance of a procedure inherits its urgency and its priority from the goal it addresses at a given instance.

Priority reflects the standard priorities involved in managing multiple trauma, which call for addressing conditions in the following order:

- airway injuries (breathing)
- circulation-related injuries (bleeding and impairments to the movement of blood)
- neurological injury
- contamination resulting from injuries
- orthopedic stability
- other conditions

Each management goal is accompanied by a priority indicator that classifies it into one of the six classes and is inherited by the procedure chosen to address it. For example, pericardial tamponade is classified as a circulation-related injury, while simple pneumothorax is classified as an airway-related injury, and thus would normally be addressed before a subclavian artery injury.

Associated with each procedure is a list of sites where it can be carried out: the Emergency Room (ER), the X-Ray Unit, the Operating Room (OR), and/or the Trauma Unit (TU). This will affect where a procedure is ordered in the plan. Also associated with each procedure is its *preconditions*, which are here limited to patient-specific contra-indications (e.g., that the patient

Rx_Simple_Pneumothorax(side=S)

type: therapeutic
urgency: nil
procedures:
tube_thoracostomy(side=S),
needle_decompression_chest(side=S)
priority: airway

Needle_Decompression_Chest(side=S)

site: ER
goals: Rx_Simple_Pneumothorax(side=S),
Rx_Tension_Pneumothorax
preconditions: NIL
priority:
urgency:

Figure 6: Procedures Knowledge Base

not have any abdominal scars from previous operations) and equipment requirements (e.g. that a CAT scan machine be available).

Figure 6 gives two examples of items in the *Procedure Knowledge Base*. The first illustrates the goal of treating a *Simple Pneumothorax* on either side. It says that this is a non-urgent therapeutic goal involving the patient's airway, that can be satisfied by either a *tube thoracostomy* or a *needle decompression* of that side of the chest. The second example illustrates the procedure *needle decompression*. It says that this is a procedure that can be performed in the Emergency Room and used for treating both a *Simple Pneumothorax* and a *Tension Pneumothorax*.

(Before going on to describe the planning algorithm, we should note that the dashed line in Figure 5 between the *Procedure Knowledge Base* and the *Reasoner* indicates a future information channel not yet incorporated into TraumAID 2.0. It will become necessary when the *Reasoner* is extended to draw diagnostic suspicions and conclusions from the relationship between the *expected effect* of an action and the patient's *current state*. Currently any such reasoning is done in an ad hoc, albeit correct, manner.)

6.2.2 Planning Algorithm

The planning algorithm carries out the following two steps:

1. It selects procedures that address the goals posted by the reasoner and orders them into a plan.

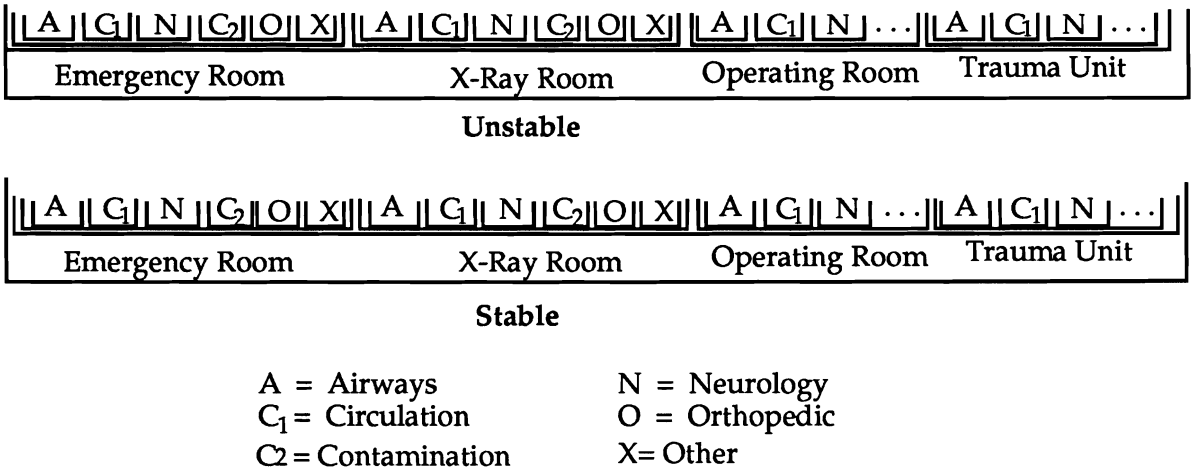

Figure 7: Basic Sorting Order

2. It then optimizes the initial procedures in this ordering to minimize risk and cost.

The planner is invoked whenever a new set of diagnostic and therapeutic goals is posted by the reasoner, using the ordering constraints that any plan must adhere to in order to drive the selection and ordering of procedures.

More specifically, input to the first planning step consists of: (1) the set of goals Γ proposed by the *reasoner*, (2) all patient-specific data acquired in response to questions and previous tests, and (3) the *Procedure Knowledge Base* described in the previous section. Given this input, the algorithm does the following:

1. It sorts the set of goals Γ based on *urgency* and *goal priority.*

2. It then constructs a plan Π through the following iterated steps, stopping when Γ is exhausted:

   (a) It picks the next goal γ on Γ. If γ is not addressed by a procedure already included in Π, it identifies the most preferred procedure π for addressing γ that does not violate patient-specific contra-indications or require equipment that is not available – i.e., the procedure that would have been chosen for that patient, were γ the only problem to be addressed.

   (b) It adds π to Π at a position that conforms to the order depicted in Figure 7. This order reflects the following precedence principles:

      i. Procedures dealing with shock and instability

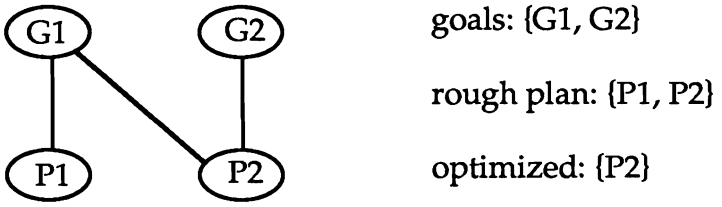

Figure 8: Remove Redundant Procedures

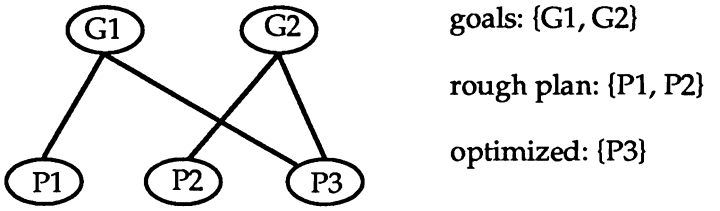

Figure 9: Prefer Broader Coverage

ii. Logistic ordering (procedures performable in the ER, followed by those performable in X-Ray, those performable in the OR, and finally those performable in the Trauma Unit)

iii. Standard practices of trauma care

iv. Therapeutic procedures for condition α take precedence over diagnostic procedures for β.

(c) It checks that π does not violate any pre-conditions or contra-indications of procedures already in the plan, with respect to orderings possible given the above constraints. If there is no way to place π in the given plan, it chooses the next best procedure that addresses γ and repeats this process.

(d) If there is no valid way of addressing γ in the current plan, it is left unaddressed, and the physician is so informed. Having ordered goals in Step 1 by urgency and priority however, any goals left unaddressed will be less urgent and less important than any goal already addressed by the plan. (We have not yet had any need to invoke this step, and hope, through prior analysis, to reduce (if not eliminate) the need to do so.)

When Γ is exhausted, the planner then acts to optimize the *initial actions* of Π. Based solely on the many-to-many mapping between goals and procedures, the planner can perform two types of optimizations: it can eliminate redundant procedures (Figure 8), and it can replace subsets of

procedures that can be *covered* by a single procedure at lower total cost and no higher total risk (Figure 9). Other types of optimizations require analyzing the *substructure* of the procedures, to identify sharable subparts and hence possible partial mergings. For example, if the first steps of a proposed plan consist of lavage and arteriogram, CAT scan might be substituted for lavage, to take advantage of the injection of dye required for the arteriogram. We have yet to do such optimizations based on common substructure . As noted, only the initial two steps of a plan are optimized, as information acquired from procedures and new information about the patient's state may lead to latter parts of the plan being delayed or superceded.

6.2.3 Brief Example

In order to focus on how the planner goes about satisfying goals set by the *rule-based reasoner*, we will only describe the activity of the planner here in detail. In the case to be discussed, an adult male presents with two knife wounds, one to the base of the *left neck*, the other to the *right lower lateral chest.* The rule-based reasoner identifies four diagnostic goals that it would like to see satisfied:

- **RO_Subclavian_Artery_Injury** - a circulation-related injury that could follow from the neck wound.
- **RO_Left_Pneumothorax** - an airway-related injury that could follow from the neck wound.
- **RO_Right_Pneumothorax** - an airway-related injury that could follow from the chest wound.
- **RO_Intra_Abdominal_Injury** - an injury associated with contamination that could follow from the chest wound.

At this point, the planner takes control. Step 1 of its algorithm considers goal urgency and priority. Because the patient is stable, none of these goals are more *urgent* than any others. By goal priority, the two airway-related goals are ordered first and second (with their relative order assigned arbitrarily), then the circulation-related goal, and finally, the contamination-related goal.

Considering the goals in this order, Step 2 uses information from the *Procedures Knowledge Base* to decide how to satisfy them. It finds that a basic chest X-ray is the procedure of choice for diagnosing a left pneumothorax. A basic chest X-ray can be performed in the Emergency Room, and is so ordered. Since it is the only procedure so far in the plan, it does not violate any pre-conditions or contra-indications of any procedures already in the plan.

Taking the next goal in the list, diagnosing a right pneumothorax is already covered by the chest X-ray procedure already in the plan, so no more need be done to the plan.

Taking the third goal, an arteriogram is the procedure of choice for diagnosing an (isolated) subclavian artery injury. Unlike a chest X-ray, an arteriogram can only be done in an X-Ray Unit, and is so ordered.

Finally, the fourth goal of diagnosing a possible intra-abdominal injury can be covered by a peritoneal lavage. Since this can be done in the Emergency Room, it is ordered before the arteriogram.

Thus the plan at the end of this first stage comprises the ordered set:

1. Chest X-Ray
2. Peritoneal Lavage
3. Arteriogram

Suppose the physician decides to perform the chest X-Ray, and finds a right pneumothorax, but no left pneumothorax. The rule-based reasoner now identifies a new goal that it would like to see satisfied (as well as the two goals still pending - RO_Subclavian_Artery_Injury and RO_Intra_Abdominal_Injury):

- **Rx_Right_Pneumothorax** - an airway-related injury caused by the chest wound.

Assuming the patient is still stable, none of these goals are again more *urgent* than any others. The first step of the algorithm thus orders the airway-related goal first, then the circulation-related goal, and finally, the contamination-related goal.

Considering the goals in this order, using information from the *Procedures Knowledge Base*, a tube thoracostomy is found to be the procedure of choice for treating a pneumothorax. It can be performed in the Emergency Room, and is so ordered.

Again, an arteriogram is the procedure of choice for diagnosing an (isolated) subclavian artery injury. An arteriogram can only be done in an X-Ray Room, and is so ordered. Again, the goal of diagnosing a possible intra-abdominal injury can be covered by a peritoneal lavage. Since this can be done in the Emergency Room, it is ordered before the arteriogram.

Thus the plan at the end of this second stage comprises the ordered set:

1. Tube Thoracostomy
2. Peritoneal Lavage
3. Arteriogram

Suppose the physician decides to perform the tube thoracostomy to treat the right pneumothorax. The rule-based reasoner now identifies a new diagnostic goal that it would like to see satisfied (again, in addition to the two goals still pending - RO_Subclavian_Artery_Injury and RO_Intra_Abdominal_Injury):

- **RO_Misplaced_Tube** - an airway-related condition associated with the insertion of the chest tube.

A follow-up chest X-ray is the procedure of choice for determining whether or not a chest tube has been inserted correctly, while nothing has changed vis-a-vis the procedure of choice for diagnosing the artery injury and intra-abdominal injury. The chest X-ray is ordered first because it involves the patient airway and can be done in the Emergency Room. Thus the plan at the end of this third stage comprises the ordered set:

1. Follow-up Chest X-ray
2. Peritoneal Lavage
3. Arteriogram

Suppose the physician performs the X-ray and finds the tube has been placed correctly. The pneumothorax is now taken to be treated, and the reasoner only posts the two diagnostic goals still pending: RO_Subclavian_Artery_Injury and RO_Intra_Abdominal_Injury.

Again, arteriogram and peritoneal lavage would be the procedures of choice for diagnosing the two injuries. On the basis of logistic ordering, the latter (which can be done in the ER) would be ordered before the former (which can only be done in an X-Ray Unit). Truly optimizing this plan requires a type of optimization the system cannot yet do. It should be able to use its knowledge that an alternative diagnostic procedure for intra-abdominal injury is a CAT scan. Since the patient can be expected to eventually have to go to the X-Ray Unit for an arteriogram and a CAT scan can take advantage of the dye that will have to be injected into the patient for the arteriogram, the system should be able to recognize that a plan comprising

1. Arteriogram
2. CAT scan

has a lower total cost than the original one. We hope to soon enable the system to perform such optimizations.

Finally, supposing the physician performs the arteriogram and finds no injury to the subclavian artery, but the CAT scan reveals liver injury. The reasoner will post, as its final goal

- **Rx_Liver_Injury**

which the planner will recommend treating with a Laparotomy.

6.3 Summary

As we hope to have made clear through the above example, the *rule-based reasoner* and the *planner* have very different roles in TraumAID 2.0 – The *rule-based reasoner* embodies the

knowledge and logical machinery needed to link clinical evidence to *goals*, and the *planner* embodies the global knowledge and logical machinery needed to address *combinations* of goals. These two roles are however linked and coordinated through the system's cycle of reasoning, planning and acting. As such, we believe the architecture of TraumAID 2.0 could well serve in other application areas as well.

7 Critiquing Physician Plans

Currently, both both versions of TraumAID act to *guide* the physician through the development and performance of a management plan. Both from our own experience and that of others, we recognize that this may not be the best way to use managment support systems in a clinical setting. In actual practice, it may be more acceptable for the system to, as Shortliffe puts it [14], "act as a sounding board for the user's own ideas, expressing agreement or suggesting reasoned alternatives." This type of behavior has been called *critiquing* [15,16,14,17]. To this end, we have augmented TraumAID 1.0 with an initial version of what we call *critique mode interaction*, in which the physician provides not only findings but also his/her proposed diagnoses and proposed management plan. The system can then critique this plan with respect to its own where appropriate, pointing out significant differences and alternative solutions. In most cases, plans will not differ significantly, and comments will be minimal. In cases where physicians are faced with injuries that the see only infrequently, the system will likely play a more active role.

A management plan in this context consists of a set of diagnostic and therapeutic procedures, the former being driven by suspicions and the latter being grouped around the diagnosis that requires them. Eventually, the aspects of a plan that TraumAID's critique-mode should address are: (1) diagnostic state reached by the physician; (2) procedures that s/he considers necessary to reach those diagnoses; and (3) the therapeutic procedures suggested to treat the patient. Currently though, TraumAID can critique only the final diagnoses and associated treatment procedures in a proposed management plan. Plans are entered through a sequence of pop-up menus, to reduce the amount of typing necessary.

7.1 Critique Features

We shall illustrate the features of TraumAID's critique mode through the following example.

> An adult male presents with a single stab wound in the midline lumbar posterior torso (mid lower back). He is not in shock but has lost motor control in his right leg as well as sensation in his left leg. X-Ray reveals no fracture of the vertabrae. The system concludes spinal cord injury and suspects abdominal wall injury. There is as yet insufficient evidence to disprove or conclude the latter.

The physician enters his or her assessment of the case, with the following possible results.

Correct partial plan. For a given diagnosis, if both the diagnosis and its associated therapeutic procedures are part of the system's plan, the system responds with approval.

Physician's plan

Diagnosis – Spinal Cord Injury
Treatment – Immobilize the patient, Neurosurgical Consultation, Paraplegia Rehabilitation.

System's response

There is enough evidence to conclude what you suggested: Spinal Cord Injury. Immobilize the patient, Neurosurgical Consultation and Paraplegia Rehabilitation are, as you mentioned, the correct protocol.

Difference in therapeutic procedures. Here, the physician's proposed therapeutic procedures for a given diagnosis do not match those of the system. The system in this case confirms the therapeutic procedures that do match and warns about those that do not. In the first of the following examples, there are no procedures on which physician and system concur. In the second, there is a single procedure that the system disagrees with.

Physician's plan

Diagnosis – Spinal Cord Injury
Treatment – Minimal Fluids, Cover Wound, Occlusive Dressing.

System's response

You have reached the diagnosis Spinal Cord Injury that matches the conclusion that the system has reached, but none of the treatments that you included in your plan, Minimal Fluids, Cover Wound and Occlusive Dressing, agrees with the standard protocol. The standard protocol is: Immobilize the patient, Neurosurgical Consultation and Paraplegia Rehabilitation.

Physician's plan

Diagnosis – Spinal Cord Injury
Treatment – Immobilize the patient, Neurosurgical Consultation, Paraplegia Rehabilitation and Minimal Fluids.

System's response

You have correctly diagnosed Spinal Cord Injury and suggested Immobilize the patient, Neurosurgical Consultation and Paraplegia Rehabilitation. However, there is no need to treat the patient with Minimal Fluids.

Missing diagnosis. It is not uncommon for further diagnostic tests to be ordered for a diagnosis already evident from the data collected from a patient. When a concludable diagnosis is missing from a physician's plan, TraumAID reports the relevant evidence, together with the role it plays in deriving the diagnosis. For teaching purposes, TraumAID also provides the set

of therapeutic procedures that address that diagnosis.

System's response (in the case where physician's plan does not include Spinal Cord Injury)

Based on the evidence:
* Wound in the Midline Posterior
* Loss sensation on the left leg
* Loss motor on the right leg
* Insufficient evidence to conclude Ischemic Spinal Cord Injury(another cause of paralysis),
it can be concluded that the patient has Spinal Cord Injury. The standard protocol is: Immobilize the patient, Neurosurgical Consultation, Paraplegia Rehabilitation.

Insufficient evidence for a diagnosis. Recall that a diagnosis may either be suspected or concluded. If the physician concludes a diagnosis that the system still only suspects (because of insufficient evidence), TraumAID suggests the physician enter its regular "let the system ask" mode that will allow it to ask for the pieces of evidence that it needs to either conclude or rule out that diagnosis. (In this example, the system makes use of the fact that it is not worth diagnosing *Abdominal Wall Injury* if the more serious injury *Non Specific Intra Abdominal Injury* has already been concluded. As long as it hasn't though, it is reasonable to pursue the lesser diagnosis.)

Physician's plan

Diagnosis – Abdominal Wall Injury
Treatment – Cover Wound Occlusive Dressing

System's response

Although there is some evidence to suspect Abdominal Wall Injury:
* Wound of type stab in the abdomen
* Not enough evidence to conclude Non Specific Intra Abdominal Injury
nothing can be concluded without more evidence. To further investigate this diagnosis, please select "let the system ask" command.

A diagnosis known to be false. If the physician includes in his/her plan a diagnosis that the system has concluded to be false, an appropiate comment is made to the physician. However since TraumAID can not determine from what it knows, the basis on which the physician reached his/her conclusion, it does not attempt to further correct the cause of the physician's possible

error.

Physician's plan

Diagnosis – Compound Fracture Vertebra
Treatment – Neurosurgical Consultation, Inmobilize the patient and Minimal Fluids.

System's response

The system has reached the conclusion that the patient does not have Compound Fracture Vertebra, because some conditions needed for that conclusion are not true.

Unmotivated diagnosis. If a diagnosis in the physician's plan is neither suspected nor concluded by the system, again the physician is informed, but is also invited to provide more information so that the system can form an opinion as well.

Physician's plan

Diagnosis – Ischemic Spinal Cord Injury
Treatment – Thoracic Aorta Repair with Shunt, Paraplegia Rehabilitation.

System's response

There is not enough evidence to either suspect or make any conclusion about Ischemic Spinal Cord Injury. To assess Ischemic Spinal Cord Injury, please provide more information.

7.2 Future Developments

We are planning several improvements to TraumAID's critique mode interaction. Most importantly, it must be extended to critique diagnostic procedures in the physician's plan which the system may find either unmotivated, in conflict with other procedures included in the plan, slower or more expensive than another one that provides essentially the same information, or unavailable at that institution. In these cases, an alternative procedure should be recommended.

Secondly, TraumAID should record dissenting opinions. We are aware that diagnostic criteria can depend on the human expert providing the system's medical knowledge (in our case, Dr. John Clarke) and are therefore subject to disagreement. Because of this, the system should be extended to record opinions from physicians when these physicians do not agree with the conclusion reached by the system. Such free-text comments could periodically be examined and used to extend the knowledge according to other experts' experience.

Thirdly, the system should optionally clarify the importance of a particular fact to the different diagnoses it may implicate. While such texts would most likely be completely "canned"

(rather than generated automatically using templates), it will undoubtly be very useful for educational applications of this critique-mode interaction.

8 Decision Support in Resource-Limited Environments

The best treatment for injuries is often available only in organized trauma centers [1]. However, not all facilities have the resources of a trauma center. One of the intriguing problems yet to be addressed systematically in expert system's research is that of accommodating the informational needs, reasoning, and recommendations of an expert system to the expertise and resources to hand. The need for such an ability is particularly apparent in multiple trauma, as there is often no choice as to where acutely injured patients are managed. The closest site may be a Level I Trauma Center or it may be a rural hospital. Automatic accommodation would eliminate the need for writing a whole new set of rules for each environment.

TraumAID represents a method of organizing the knowledge that produces superior results into an aid that can be disseminated to appropriate local and regional facilities for the acute care of injured patients presented to them. We have already demonstrated that TraumAID 1.0 can be adapted for use in health care settings with fewer resources and personnel with less expertise, such as for use by independent-duty medical corpsmen onboard submarines [18]. Instead of creating a new rule base, we modified the version of the system designed for use by physicians in a well-equipped Trauma Center. The modifications retain the same rule base but accommodate the absence of particular diagnostic tools (e.g., radiology) through a set of *safe assumptions* represented as default values. A safe assumption is used if its corresponding test cannot be performed because of a lack of equipment or training. Conclusions reached in this way are identified as such.

For example, TraumAID will suspect bladder injury in cases where a patient has sustained a wound in his lower abdomen and shows signs of hematuria. Conclusive evidence would come with a cystogram. However, this test is unavailable to submarine corpsmen, so TraumAID makes a safe assumption – in this case, to act as if the result were *positive* – which leads it to recommend that the patient be evacuated urgently for possible bladder repair with chromic sutures and drainage.

To accommodate the corpsmen's setting and skills, TraumAID's therapeutic recommendations were systematically translated to ones appropriate for the corpsmen, most often involving observation or evacuation. After this systematic translation, it was only necessary to hand-modify one set of rules (9 rules in all) dealing with the diagnosis of hemopneumothoraces, in order for the system to conform to its new environment. These modifications all involved a test

so fundamental – a chest x-ray – that no standard default was possible. Instead, another test (auscultation of the lungs) was substituted which, although less reliable, can be performed by a corpsman in an environment that lacks x-ray resources.

We have yet to consider the problem of resource adjustment in more generality, but it is clearly a useful direction in which expert systems can develop.

9 Conclusion

The ultimate goal of TraumAID, like that of other medical decision support systems, is higher quality care at lower cost. In this light, the most significant potential use of TraumAID is to provide physicians with *timely* recommendations for the initial definitive management of injuries – recommendations which would be comparable to the advice of experts. As patient protocols for individual patients, they would be consistent with protocols that minimize the chance of errors leading to avoidable mortality, morbidity or disability. These protocols would form standards of care that could also be used for *quality assurance* as care was being given. TraumAID is potentially a broad-based solution to controlling injury mortality, morbidity and disability that result from sub-optimal management plans during initial definitive hospital care. Even at trauma centers, it would be a way of monitoring standards and protocols for both care and studies.

Over five years of effort has gone into this work on TraumAID, and we can easily see another five years as we validate what we have done, broaden our coverage, and replace limited, *ad hoc* methods of reasoning with cleaner, more extensible ones. Working on TraumAID has been both an education and a pleasure for all members of the group, and we believe that the results are a contribution to both trauma management and Artificial Intelligence.

References

[1] Committee on Trauma Research. *Injury in America.* Washington: National Academy Press, 1985.

[2] Trunkey, D.D., Trauma, *Scientific American*, 249, pp. 28-35, 1983.

[3] West, J.G., Trunkey, D.D. and Lim, R.C., Systems of Trauma Care: A Study of Two Counties. *Archives of Surgery*, 114, pp. 455-460, 1979.

[4] American College of Surgeons, *Advanced Trauma Life Support.* Chicago: American College of Surgeons,1984.

[5] Moore, E.E., Eiseman, B. and Van Way, C.W., *Critical Decisions in Trauma*, St. Louis: CV Mosby, 1984.

[6] Champion, H.R., Sacco, W.J., Carnazzo, A.J. et al., Trauma Score, *Critical Care Medicine*, 9, pp. 672-676, 1981.

[7] Teasdale, G. and Jennett, B. Assessment of Coma and Impaired Consciousness: A practical scale. *Lancet* 2, 1974, pp.81-84.

[8] Clarke, J.R., Cebula, D. and Webber, B. A Computerized Decision Aid for Trauma, *Journal of Trauma* 28, 1988, pp.1250-1254.

[9] Clarke, J.R., Niv, M. and Webber, B.L. Computerized Patient-specific Protocols for Managing Penetrating Chest Injuries (abstract). *Theoretical Surgery* 5(3), 1990, p.148-149.

[10] Buchanan, B. and Shortliffe, E., *Rule-Based Expert Systems.* Boston: Addison-Wesley,1984.

[11] Poole, D. and Provan, G. What is an Optimal Diagnosis? *AAAI Workshop on Reasoning under Uncertainty*, Boston MA, July 1990.

[12] Rymon, R., Webber, B., and Clarke, J.R. Responsive Planning in TraumAid. *Proc. 6th Israeli Conference on Artificial Intelligence and Computer Vision*, pp. 257-271, December 1989.

[13] Rymon, R., Webber, B., and Clarke, J.R. Progressive Horizon Planning. Submitted to the *7th Israeli Conference on Artificial Intelligence and Computer Vision*, December 1990.

[14] Shortliffe, E. Decision Support Systems. In E. Shortliffe et al. (eds.) *Medical Informatics: Computer Applications in Health Care.* Reading MA: Addison-Wesley, 1990.

[15] Langlotz, C. and Shortliffe, E. Adapting a consultation system to critique user plans. *Intl. J. Man-Machine Studies* 19, pp.479-496, 1983.

[16] Miller, P.L. Attending: Critiquing a physician's management plan. *IEEE PAMI* 5, pp. 449-461, 1983.

[17] van der Lei, J. and Musen, M. Critiquing Physician Decision Making using Data from Automated Medical Records: an evaluation of HYPERCRITIC. *AAAI Spring Symposium on AI in Medicine*, Stanford University, March 1990.

[18] Clarke, J.R., Niv, M., Webber, B., Ryack, B., et al. TraumAID: A decision aid for managing trauma at various levels of resources. *Proc. 13th Annual Symposium on Computer Applications in Medical Care*, Washington D.C., November 1989, pp. 1005-1006.