HEURISTIC MODELING USING RULE-BASED COMPUTER SYSTEMS

D. A. Waterman and Brian M. Jenkins

March 1977

P-5811

The Rand Paper Series

Papers are issued by The Rand Corporation as a service to its professional staff. Their purpose is to facilitate the exchange of ideas among those who share the author's research interests; Papers are not reports prepared in fulfillment of Rand's contracts or grants. Views expressed in a Paper are the author's own, and are not necessarily shared by Rand or its research sponsors.

The Rand Corporation
Santa Monica, California 90406

ABSTRACT

This Paper discusses a recent effort at The Rand Corporation to apply heuristic modeling techniques to the topic of international terrorism. By heuristic model we mean an information processing model described in terms of heuristics (rules of thumb) which define the operation of the model. This joint project by computer scientists and social scientists is designed to provide researchers working on the problem of terrorism with a tool, in the form of computer programs, which will aid them in their analysis of terrorist activities. As a first attempt at developing such a tool, we have precisely defined the basic concepts needed for a model of terrorism and have used these concepts in the implementation of a program that extracts information about terrorist activities from a terrorism analyst and automatically represents it in a form that facilitates machine analysis of the information. This analysis consists of applying rules to the data extracted from the analyst in order to infer new data. We have found that the exercise of extracting information from an analyst in terms of precisely defined concepts not only sharpens the analyst's skills but also provides a basis for constructing programs to process the information.

I. INTRODUCTION

This Paper describes work in progress on the task of applying heuristic modeling techniques to ill-defined problems in the social sciences. The ultimate goal of the project is to explore ways in which current computer technology, in the form of rule-based systems, can be used to help an expert understand the domain he is working in and assist him in decisionmaking, particularly in crisis situations. Our approach is to provide the expert with a set of computer programs that constitute a model of his domain of expertise and that enable him to store information explicitly as rules and data. He can then use the model as a tool that guides and stimulates decisionmaking by its ability to explain the lines of reasoning it uses to arrive at each decision it makes.

The heuristic model is a model of a situation stated in terms of heuristics (rules of thumb) which describe the dynamics of the situation. It is particularly useful for problem domains that are not well formalized and for which no generally agreed upon axioms or theorems exist. The domain we have chosen for investigation -- international terrorism -- is just this type of domain. By studying the techniques needed to develop a heuristic model of international terrorism we hope to gain insight into how these methods can be applied to other ill-defined domains in the social sciences.

In this study we outline the objectives of the project, describe how information can be represented in rule-based form, describe some preliminary attempts to create programs that help

the expert input information into the computer system, and finally, discuss the usefulness of providing an expert with this type of tool.

BACKGROUND

This project is defined by the intersection of two separate research efforts at Rand -- one involving the application of computer science techniques to the problem of helping people who are not computer experts interact with computers, and the other involving the study and analysis of international terrorism. The immediate goal is to provide terrorism analysts with computer programs specifically designed to aid them in their analysis of terrorist activities. These programs are designed so that they can be carefully tailored to the individual needs of each analyst.

The computer science project, initiated in July of 1974, is based on the premise that a good way to aid a computer user is to provide him with his own personal computer, one that is sensitive to his needs and desires. This can be accomplished in many ways (via timesharing, intelligent terminals, etc.) but the underlying theme is to have the software designed to accommodate the skills of the user, with more sophisticated software for the less advanced users. The system developed for this is called RITA, Rule-directed Interactive Transaction Agent (Anderson and Gillogly, 1976a, 1976b). An interactive transaction agent is simply a program (piece of software) that can perform some task for the user. Examples of typical tasks are (a) filing, retrieving, and editing of data on local storage files, (b) handling interactive dialogues with external information systems,

(c) providing a local tutorial facility, and (d) heuristic modeling of subjective or judgmental information (Waterman, forthcoming). The most distinctive feature of the RITA architecture is the use of rules to describe judgments, where a rule is defined to have the form "IF condition THEN action," meaning "if the given condition is true in the current situation then perform the recommended action." This is called a rule-based or production system approach to heuristic modeling.

Since 1973, The Rand Corporation, sponsored jointly by the Defense Advanced Research Projects Agency and the Department of State, has been engaged in research on the phenomena of international terrorism. The theory of terrorism, terrorist tactics, and terrorist groups has been examined, and data on over a thousand incidents of political violence have been amassed. Several hundred of these have been the subject of more detailed examinations. During the course of the research, the participants in the project came naturally to make "expert judgments" or to develop "hunches" regarding various aspects of terrorist activity, such as the responsibility for incidents where credit was not claimed, the objectives of a particular action, and the probable outcome of certain types of episodes. Though often unarticulated, and sometimes unrealized by the participant, these hunches were often based on a series of individual "rules" derived from an understanding of the logic behind the use of terrorist tactics, the observed modi operandi of various terrorist groups, and the outcome of similar episodes in the past. Thus the domain of terrorism is a reasonable area in which to use rule-directed systems.

The heuristic modeling project is beneficial to both the

computer scientists involved in the intelligent terminal project and the social scientists working on the terrorism analysis project. The computer scientists benefit by gaining access to a problem domain which can be used to evoke and test ideas about the design and implementation of heuristic models. The feasibility of using a rule-based system like RITA for constructing a large complex model of decisionmaking in an ill-defined domain has already been demonstrated by the MYCIN system (Shortliffe, 1974; Davis, et al., 1975). MYCIN is a computer model composed of over 200 rules which provide the basis for medical consultations on the task of selecting antibiotic therapy for patients with bacterial infections.

Heuristic modeling is also beneficial to the analysts studying terrorism. It provides a requirement and a means for articulating the series of steps that lie behind the "intuitive" judgments made by the analysts in reaching a particular conclusion. In the process of explaining each step, they find themselves asking questions not asked before, simply because the rapid process of thought does not demand such detailed explanations. Previously, in describing why they felt something to be so, it was explained simply as an intuitive judgment, or a hunch. The requirements of machine analysis impose a new degree of rigor on the analysts, lead to new questions, extract articulations about the theory and logic of terrorism, and point to new areas of exploration. In sum, it not only provides a useful tool for the analysts to use but also aids them in formalizing their domain of study.

When fully operational the heuristic model or "intelligent agent" will continue to demand of the analyst that

he examine each incident in a comprehensive manner, and that he explain each judgment. On the basis of rules provided by the analyst the agent will reach conclusions with which the analyst may disagree. In this case, the analyst will be compelled to reexamine the basis for his own deduction, and this may lead to the formulation of new rules for the agent. In this way the system will enable analysts to preserve their own judgments and those of others in a form that is easily retrievable and can be made available to decisionmakers faced with real crisis situations.

BASIC ASSUMPTIONS

This project is based on three major assumptions. The first is that terrorist activities can be analyzed in a formal, rigorous fashion. The violence of political extremists, for the most part, is not viewed as a collection of mindless, irrational, and thus totally unpredictable events. It follows certain rules of logic that can be discerned by the experienced analyst. Patterns emerge which can be translated into "rules." For example, a study of political kidnappings and hostage incidents shows that political extremist groups that operate in their home territory and have the support of an underground organization generally prefer standard kidnappings as a means of taking hostages. (In a standard kidnapping, the victim is held at an unknown location while the kidnappers bargain by means of telephone or mail with the targets of their demands. This enables them to hold their hostage for months if necessary to increase pressure for concessions.) Groups operating abroad or

lacking an underground tend to become involved in barricade-and-hostage incidents (captors seize their hostages in a public place allowing themselves to become hostages; they bargain for escape along with other demands). This is not always true, but it is generally true. Some of the patterns are based on group characteristics or modi operandi. For example, the Irish Republican Army has shown no inclination to seize hostages for bargaining purposes. Many of the Palestinian and Latin American groups frequently do so.

The second assumption is that although the domain of terrorism is obviously complex with each group and each incident highly unique, it is nonetheless believed that even a limited number of rules that prod thinking and formalize the domain will enhance analysis of the topic.

The third assumption is that heuristic models of terrorist activities can be formulated as rule-directed models. We have already justified the use of the rule-based architecture through reference to the work done with MYCIN. The limiting constraint here is the number of rules needed to adequately model the domain. This can be avoided to a certain extent by attempting to model only a small portion of the domain of international terrorism. In this initial effort we have restricted our attention to the problem of terrorist bombings.

OBJECTIVES OF THE PROJECT

The primary objective of the heuristic modeling project is to determine the feasibility and utility of developing heuristic models within the rule-directed framework and to help

the user make complex decisions in ill-specified domains containing large amounts of information that must be coordinated and used in a short amount of time. The secondary objective is to develop a demonstration model that performs interesting and useful deductions within the domain of international terrorism. Other objectives are:

(1) to provide a method or procedure to stimulate new ideas and provoke new questions about the domain.

(2) to provide a framework within which a high degree of analytic rigor can be imposed.

(3) to provide a framework for decisionmaking that structures thinking and guides problem solving during a crisis situation.

(4) to assist in the analysis and understanding of terrorist activities and groups.

(5) to ultimately provide a predictive tool, one that can infer what will happen in the near future or what the immediate effect of certain actions will be.

CHARACTERISTICS OF THE HEURISTIC MODEL

The heuristic model currently under development is a collection of RITA agents that the terrorism analyst will have at his disposal. These agents will constitute a tool that the analyst can use to aid in decisionmaking or problemsolving tasks. The model will have the following primary characteristics:

(1) A deductive inference capability. This means that the agents will be capable of not only retrieving information stored in the data base but of using that information to make deductions that can cause new information to be added to the data

base. This capability is a direct result of using a rule-based model, since the rules are, in effect, procedures defining how new information can be deduced from existing data.

(2) Dynamic growth. In the domain of international terrorism both data and judgments (or policy) are subject to constant revision. The modularity imposed by the rule-based organization of the agents facilitates such dynamic growth, i.e., it is easy to add, delete, or modify both the data and the rules that constitute the judgments.

(3) Integrated collection of knowledge. The agents that constitute the model can represent judgments of many experts from various parts of the world. This expertise is stored at one location for later use by experts and nonexperts alike. However, the real utility of such a system is that it will permit a constant exchange of data and judgments by leading experts, and this will lead to the formation of new rules and ideas about the domain.

II. REPRESENTATION OF INFORMATION IN PRODUCTION SYSTEM FORM

A production system is a collection of rules of the form condition --> action (Newell and Simon, 1972), where the conditions are statements about the contents of a data base and the actions are procedures that alter the contents of that data base. RITA is a specialized production system architecture, that is, a programming language within which one can write specific programs that are production systems or RITA agents. Production system architectures have been developed to facilitate adaptive behavior (Waterman, 1970, 1975; Waterman and Newell, 1976), to model human cognition and memory (Newell, 1972, 1973), and to create large systems that represent heuristic models of the judgments of a collection of experts (Feigenbaum, et al., 1970; Shortliffe and Buchanan, 1975; Davis, et al., 1975; Lenat, 1976). RITA was designed with the latter task in mind and accordingly has two types of production rules: those called rules and those called goals. The rules are left-hand-side driven, i.e., when all the conditions of a rule are true relative to the data base, the rule "fires," causing the associated actions to be taken. The goals, however, are right-hand-side driven. This means that the system is given a condition to make true, or, in effect, a question to answer through deductive inference. The right-hand sides of rules are examined to find one that could make the desired condition true. When such a rule is found, its left-hand side is examined to see if all its conditions are true. If they are, the rule is fired; if not, the process continues in the same manner in an attempt to make each condition in the left-hand side of the rule true.

DATA REPRESENTATION IN RITA

In RITA the data are represented as objects that can have any number of attribute-value pairs. Thus to represent a person whose name is John Smith, whose age is 32, and whose salary is in the $33,789 to $43,923 range, the object PERSON would have associated with it three attribute-value pairs: NAME = JOHN SMITH, AGE = 32, and SALARY RANGE = $33,789 TO $43,923. This would be written in RITA as:

```
OBJECT person
    name is "John Smith",
    age  is "32",
    salary-range is "$33,789 to $43,923";
```

The fact that John Smith, Mary Jones, and Tom Brown are all government service employees and are all part of a group of such employees called GS level 15, can be represented as:

```
OBJECT group
    name is "GS Level 15",
    type is "government service employees",
    members is ("John Smith","Mary Jones","Tom Brown");
```

Note that the value of an attribute can be either an item like "John Smith" or a list like ("John Smith", "Mary Jones", "Tom Brown"). Also, more than one object of the same type can exist simultaneously in the data base. The data base can contain many distinct persons and groups, each with different attribute-value pairs. Thus Mary Jones might also be in the data base as:

```
OBJECT person
    name is "Mary Jones",
    age is "22";
```

GOAL AND RULE REPRESENTATION IN RITA

Judgments and procedures in RITA are represented as goals and rules. The goals are more generally used for judgments and the rules for procedures, but either can be used for judgments or procedures. The format for a goal is shown below.

```
GOAL goalname
        IF    premise AND premise ... AND premise
        THEN  action AND action ... AND action.
```

A few examples will make this format clear. Figure 1 shows judgments in English and below them the corresponding RITA goals.

1. "The salary range of a government service group can be determined by the salary range of any of its members."

or

"If you don't know the salary range of a government service group but you do know the salary range of a member of that group then the salary range of the group is just the salary range of that member."

```
GOAL 1

    IF   there is a group whose type is
            "government service employees"
         and whose salary-range is not known
         and there is a person whose salary-range
            is known
         and the name of the person is in the members
            of the group

  THEN   set salary-range of the group to the
         salary-range of the person;
```

2. "The salary range of a member of a government service group can be determined by the salary range of the group."

or

"If you don't know the salary range of a member

of a government service group but the salary range of the group is known then the salary range of the person is just the salary range of the group."

GOAL 2

IF there is a person whose salary-range is not known and there is a group whose salary-range is known and the name of the person is in the members of the group

THEN set the salary-range of the person to the salary-range of the group;

Fig. 1 -- Examples of RITA goals

The format for rules is very similar to that for goals as shown below.

RULE rulename

IF premise AND premise ... AND premise

THEN action AND action ... AND action.

Some examples of RITA rules are given in Fig. 2. These tend to describe processes or procedures more than judgments.

1. "Send the name and salary range of each person to the user."

or

"If the salary range of a person is known and it has not yet been sent to the user then send both the name and salary range of the person to the user".

RULE 1

IF there is a person whose salary-range is known and whose status is not known

THEN send the name of the person to the user and send the salary-range of the person to the user and set the status of the person to "accounted for";

2. "Infer the salary range of every person whose salary range is not known."

or

"If there is a person whose salary range is not known then infer that salary range."

```
RULE 2

   IF   there is a person whose salary-range is not known

 THEN   deduce the salary-range of the person;
```

Fig. 2 -- Examples of RITA rules

The "deduce" action in rule 2 is a signal to RITA to process the goals in an attempt to infer the desired information.

The data, goals, and rules just described can be combined to form a RITA agent which will tell the user the name and salary range of each person in the data base if that range is known or can be deduced. The agent is shown in Fig. 3.

```
OBJECT person<1>:
      name     IS   "John Smith",
      age      IS   "32",
      salary-range is "$33,789 to $43,923";

OBJECT person<2>:
      name     IS   "Mary Jones",
      age      IS   "22";

OBJECT group<1>:
      name     IS   "GS Level 15",
      type     IS   "government service employees",
      members IS    ( "John Smith", "Mary Jones", "Tom Brown" );
```

```
RULE 1:
   IF: THERE IS a person WHOSE salary-range IS KNOWN
               AND WHOSE status IS NOT KNOWN
 THEN: SEND the name OF the person TO user
       & SEND the salary-range OF the person TO user
       & SET the status OF the person TO "accounted for";


RULE 2:
   IF: THERE IS a person WHOSE salary-range IS NOT KNOWN
 THEN: DEDUCE the salary-range OF the person;


GOAL 1:
   IF: THERE IS a group
            WHOSE type IS "government service employees"
            AND WHOSE salary-range IS NOT KNOWN
       & THERE IS a person WHOSE salary-range IS KNOWN
       & the name OF the person IS IN the members OF the group
    THEN: SET the salary-range OF the group TO the
            salary-range OF the person;


GOAL 2:
   IF: THERE IS a person WHOSE salary-range IS NOT KNOWN
       & THERE IS a group WHOSE salary-range IS KNOWN
       & the name OF the person IS IN the members OF the group
 THEN: SET the salary-range OF the person TO the
         salary-range OF the group;
```

Fig. 3 -- A RITA agent that attempts to deduce the salary range of all persons in the data base

EXAMPLE OF A DEDUCTION BY A RITA AGENT

When the agent in Fig. 3 is executed, rule 1 is tested against the data base and found to be true, since John Smith's salary range is known but nothing is known about John Smith's status. Since all the premises of rule 1 are true, rule 1 "fires", i.e., the associated action is taken and the agent writes "John Smith $33,789 to $43,923" at the user's terminal and sets John Smith's status to "accounted for." Again rule 1 is tested against the data base but now fails to fire for John Smith, since the second premise, concerning John Smith's status,

is no longer true. It also fails to fire for Mary Jones, since the first premise concerning her salary range is not true.

Since the premises in rule 1 are not true, it can not fire and rule 2 is tested. The premise in rule 2 is true for Mary Jones so the rule fires and the salary range of Mary Jones is deduced. This deduction involves just the goals. The right-hand sides (action parts) of the goals are checked to see if they could determine the salary range of a person. Only goal 2 can do this so its premises are checked to see if they are true so the goal can be fired.

The first premise of goal 2 is true for Mary Jones but the second premise is not, since the salary range of the group is not known. Now the system does a clever thing. Instead of giving up and deciding that goal 2 is not true it tries to deduce the second premise of goal 2, i.e., the salary range of the group. The only applicable goal is goal 1, so its premises are checked. The first premise of goal 1 is true, since the group is government service employees and its affiliation is not known. The second and third premises of goal 1 are true for John Smith and the group, so all the premises are true and the goal fires, setting the salary range of the group to "$33,789 to $43,923." This makes the second premise of goal 2 true, and, since the third premise of goal 2 is true for Mary Jones, all the premises are true and goal 2 fires, setting the salary range of Mary Jones to "$33,789 to $43,923."

The deduction is now complete and control returns to the rules. Rule 1 is checked and this time is true for Mary Jones. It fires, printing "Mary Jones $33,789 to $43,923" at the user's

terminal and setting the status of Mary Jones to "accounted for." Now rule 1 is checked again but is not true, nor is rule 2, so the agent halts. The data base after the deductions are made is shown below.

```
OBJECT person<1>:
        name     IS   "John Smith",
        age      IS   "32",
        salary-range     IS   "$33,789 to $43,923",
        status   IS   "accounted for";

OBJECT person<2>:
        name     IS   "Mary Jones",
        age      IS   "22",
        salary-range     IS   "$33,789 to $43,923",
        status   IS   "accounted for";

OBJECT group<1>:
        name     IS   "GS Level 15",
        type     IS   "government service employees",
        members IS    ("John Smith","Mary Jones","Tom Brown"),
        salary-range     IS   "$33,789 to $43,923";
```

Note that not only was the desired information about the salary range of Mary Jones deduced and added to the data base but as a side effect information concerning the salary range of the group was also deduced and stored in the data base.

III. HEURISTIC MODELING USING RITA AGENTS

Building a heuristic model for a domain like terrorism is a complex task independent of the method chosen for constructing the model. The method used here centers on extracting data and rules from an expert on terrorism in a form that focuses on the essential and relevant information in the domain.

DATA AND RULE EXTRACTION

The first step involved in building a heuristic model in an ill-defined domain is the job of formalizing the domain. This critical step consists of deciding what elements in the domain are relevant to the problem at hand, how they should be categorized, and how they should be defined. It imposes the degree of organization on the data that is needed for later development of relationships between the data elements. The key to formalizing the domain is data extraction -- the process of extracting data or knowledge about the domain from an expert. This was accomplished, in our case, by a series of dialogues between the expert and the model-builders, in which the model-builders played the role of protocol analysts (Waterman and Newell, 1973, 1976). Typically the protocol analyst will ask pertinent questions of the expert which reveal assumptions and attitudes the expert may not have been conscious of having. These ideas are then expanded by the expert and the protocol analyst into a characterization of the domain.

Rule extraction proceeds in exactly the same manner as

data extraction, through an extended dialogue between the expert and the protocol analysts. We have found that trying to extract rules out of context, i.e., posing the question, "Give me all the rules you know about bombings by the Palestinian Groups," leads to vague generalities that are seldom useful. The most successful approach has been to focus on particular events, analyze them in great depth, and then classify the rules generated during the analysis as either specific to a particular context or as a special case that is generally true -- one that almost occupies the status of an axiom or tautology in the system. We thus end up with two types of rules: axiomatic ones that are true with near 100 percent probability, and heuristic ones that are true with a lesser probability. While both types of rules can be used to infer new information, the axiomatic rules can also be used to check the validity of the information being extracted.

Data extraction or domain formalization is not something that is done in a single step that is followed by the step of rule extraction. In actuality these steps proceed in parallel, progress in each affecting the status of the other. Thus many iterations must be made through the cycle of defining and organizing relevent data, generating rules that relate data elements to each other, and then revising the definitions and organization to accommodate new components of new rules. Eventually we would like to have the job of rule extraction automated, performed perhaps by a RITA agent capable of asking the right questions and using the answers to build RITA rules.

FOCAL POINTS FOR TERRORISM

It appears to us that there are three important focal points around which we can build a system for modeling terrorist activities. These are focal points around which information can be gathered and rules can be formulated. They are "event," "group," and "context." The term "event" stands for a terrorist incident: a bombing, hijacking, kidnapping, etc. The term "group" means a terrorist group, that is, a group that has used terrorist tactics. "Context" is a somewhat less precise label that encompasses local political and economic developments, effectiveness of internal security, etc. In a specific situation, it may include the circumstances external to the actual incident or event, i.e. the response of a government to the specific terrorist activity.

A useful analogy is to consider these focal points as hubs of wheels; one wheel for group, one for event, and another for context. Radiating from the hub like spokes are the "attributes" that are associated with each wheel. Radiating from the hub labeled "event" are attributes pertaining to its date, time, type, location, target, victim, and so on. Protruding from the hub labeled "group" are attributes such as size, ideology, composition of membership, etc. Around "context" would be attributes like the effectiveness of police, etc. The answer or item of information associated with each attribute, the "value," is located at the rim of the wheel (see Fig. 4).

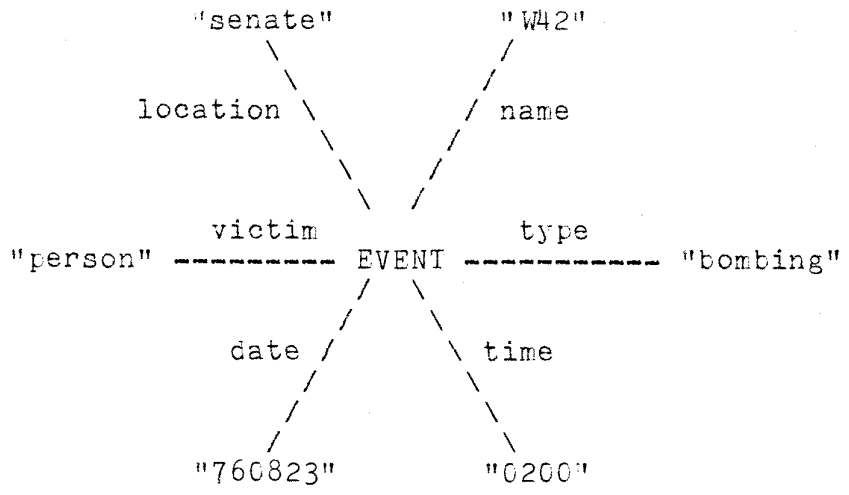

Fig. 4 -- Example of object-attribute-value relationship for the object "event"

Rules can relate one attribute or spoke to another on the same focal point or wheel. For example, the answer (or value) to a specific question (that derives from an attribute) will determine the answer to another question. To illustrate, let us use a bombing (an event). One of its attributes is the "target." By stating the target (specific individual, political property, corporate headquarters or offices, etc.) the machine may be led to a deduction about the tactical objective of the event. For example, if the target is "vital systems," the tactical objective may be "disruption."

The rules may also nullify certain questions on the basis of answers to previous questions. For example, if the bomb is identified as a "Molotov cocktail," the machine does not need to ask how the bomb was delivered. Molotov cocktails are thrown. It does not need to ask about the size of the bomb, since it knows it must have been small. These rules do not represent sophisticated judgments for which experts are required. They are

virtually axiomatic. They are included only to expedite the interrogation process by dropping irrelevant or repetitious questions.

Rules may also relate attributes from one focal point to another. For example, if we are talking about bombings, the type of bomb and type of target regularly chosen by a particular group may tell us something about the educational background of its members or about its ideology. Such rules are not likely to be axiomatic. The judgments here are likely to be quite complex and based upon several or even many answers to specific questions.

The possible combinations are endless. Rules may be built upon several attributes in each focal point. To arrive at a judgment of what members of the Japanese Red Army who are holding hostages in the French Embassy are most likely to do if its demands are not met requires a very complex chain of reasoning. However, it is important to point out that the deduction does not have to be correct so long as the user can follow the chain of logic used by the machine to arrive at its conclusion. The user then can agree with certain parts or disagree to reach his own conclusion. In the process, he will have been compelled to review his own reasoning and examine his own assumptions. The machine will have assisted, not replaced, the analyst or decision-maker.

IV. CURRENT STATE OF DEVELOPMENT

At this point we have made one iteration through the cycle of formalizing the domain of terrorist bombings and refining the associated rules. Table 1 shows a small representative sample of the information that represents our concept of the domain. The entire formalization is given in Appendix A.

```
event  ............

   name
   [give the event to be defined an identifying name]

   type
   [type of the terrorist event; e.g. bombing, kidnapping]
      bombing
      barricade-and-hostage
      kidnapping

   date
   [use the format yymmdd; e.g., 760330]

   target
   [who or what was the immediate object of the attack?]
      specific individual
      representative individual
      random individual
      political property
      corporate headquarters or offices
      other commercial property
      private property
      vital systems
      head of state

   instigator
   [name of person(s) or group who planned the event]

   perpetrator
   [name of person(s) or group carrying out the event]


group  ............

   name
   [type the full proper name used by the group and
    the acronym, if any]

   objective
```

```
   [what are the strategic objectives of the group?]
      enlarge the group
      expansion
      overthrow of government
      discredit the system
      territorial independence
      limited reforms or concessions
      vigilante-like defense

   tactics
   [the typical types of events perpetrated by the group]
      bombing
      barricade-hostage
      kidnapping
      assassinations
      highjackings
      armed assaults
      robbery
      dissemination of propaganda
      sabotage


person  ............

   name
   [name of the person]

   sex
   [male or female]

   age
   [approximate age in years]

   group-membership
   [organizational affiliation]
```

Table 1 -- Sample of Formalization of Terrorism Domain for Terrorist Bombing Incidents (see Appendix A)

To aid the terrorism expert in construction of the domain, a RITA agent called MODIFY has been implemented. This agent queries the expert, permitting him to add, modify , or delete information from the data base that represents the formalization of the domain. The expert uses MODIFY first to create the data base and then to update and modify it as his conception of the domain becomes more clear. A protocol of an expert interacting with the MODIFY agent to add a new attribute

called STAGE to the object GROUP is shown below. The expert's answers are given in italics.

What is to be modified...*group*

What attribute of group is to be modified...*stage*

Now defining the stage of the group:

What is the type number? *?*
Let me rephrase the question...

Do you want to enumerate some of the stages of the group? *yes*
Can any single group have more than one stage? *no*

Give a prompt for the stage of the group
the current stage of the development of the group (1 through 5)

Now enter new values and prompts
[Type only a carriage return to terminate]

A new value for the stage of the group is...*1*
The prompt for "1" is...
violent propaganda stage: sporadic bombings of symbolic targets

A new value for the stage of the group is...*2*
The prompt for "2" is...
organizational growth stage: attacks with specific tactical objectives

A new value for the stage of the group is...*3*
The prompt for "3" is...
guerilla offensive stage: attacks on real rather than symbolic targets

A new value for the stage of the group is...*4*
The prompt for "4" is...
mobilization of masses stage: attack on judiciary to provoke repression

A new value for the stage of the group is...*5*
The prompt for "5" is...
urban uprising stage: mass uprising, full scale urban warfare

A new value for the stage of the group is...

What attribute of group is to be modified...

What is to be modified...
Objects saved on file bomb.objects

After the expert has created the data base that represents the domain he has at his disposal, another RITA agent, called BUILD, uses the domain just defined to help him enter new data about terrorist events or groups into the computer. BUILD uses the domain data base to formulate questions about the object to be described. For example, the questions BUILD would ask as a result of the definition given to the MODIFY agent are shown below in an actual protocol of an expert interacting with BUILD. Again the expert's responses are shown in italics.

```
Is the stage of the group :
  1? ?

[violent propaganda stage: sporadic bombings
 of symbolic targets]

  1? no
  2? ?

[organizational growth stage: attacks with
 specific tactical objectives]

  2? no
  3? ?

[guerrilla offensive stage: attacks on real
 rather than symbolic targets]

  3? no
  4? ?

[mobilization of masses stage: attack on
 judiciary to provoke repression]

  4? no
  5? ?

[urban uprising stage: mass uprising, full
 scale urban warfare]

  5? yes
```

In this protocol the expert responded to each question of the agent by first typing a question mark to obtain a definition of the term being used and then answering with either "yes" or "no."

This feature of being able to ask the agent to explain itself is of critical importance, since the definitions of terms are somewhat arbitrary and not easily inferred. After the expert answered "yes" to the last question, BUILD created the object shown below.

```
OBJECT group<1>:
       stage   IS   "5";
```

The BUILD agent is designed to help the user enter information about terrorist activities into the computer in a form that can be accessed by other RITA agents. However, BUILD does more than just query the user and map the user's replies into RITA data. Before it asks each question it attempts to deduce the answer itself based on the information it has gathered up to that point. If it can deduce the answer it notifies the user of that fact and does not ask him the question. Only when it is unable to deduce the answer does it query the user. Some typical RITA goals that BUILD uses to make deductions about bombing incidents are shown in Fig. 5. These goals have the status of axioms in the system, i.e., the probability of their being true is very close to 1.

```
GOAL 1:
    IF: the time-of-day OF the event IS less THAN 0600
        & the bomb-size OF the event IS "small"
  THEN: PUT "property damage-no casualties" INTO the
          tactical-objective OF the event AS the LAST MEMBER;

GOAL 2:
    IF: "no warning" IS IN the recipient-of-warning OF the event
  THEN: SET the warning-time OF the event TO 0;

GOAL 3:
    IF: the bomb-type OF the event IS "Molotov-cocktail"
  THEN: SET the bomb-delivery OF the event TO "thrown"
        & SET the bomb-detonation OF the event TO "impact"
```

```
      & SET the secondary-detonation OF the event TO "no detonation"
      & SET the bomb-sophistication OF the event TO "low"
      & SET the detonation-source OF the event TO "self-detonating"
      & PUT "home-made" INTO the explosive-source OF the event
        AS the LAST MEMBER
      & SET the bomb-size OF the event TO "small";


GOAL 4:
    IF: the bomb-type OF the event IS "letter or parcel"
  THEN: SET the bomb-delivery OF the event TO "delivered"
        & SET the bomb-size OF the event TO "small"
        & SET the risk-to-perpetrator OF the event TO "low";


GOAL 5:
    IF: the bomb-type OF the event IS "land mine"
  THEN: SET the bomb-delivery OF the event TO "planted"
        & PUT "no warning" INTO the recipient-of-warning
          OF the event AS the LAST MEMBER;
```

Fig. 5 -- Some of the Goals used by the BUILD agent to deduce information about terrorist bombings

A complete protocol of an expert entering data about a hypothetical bombing incident is shown in Appendix B. An informal description of the hypothetical incident plus the resulting RITA description of it are shown in Fig. 6.

> The event took place in Washington D.C. at about 1:26 pm in the Senate Office Building on August 23, 1976. As Secretary of State Henry Kissinger was exiting the front door a lone individual hurled a Molotov-cocktail at him and then fled on foot. The bomb failed to explode and no one was hurt. Two minutes before the attempted bombing a woman who identified herself as a member of the New World Liberation Front called the Washington Post and claimed credit for the act. She said the reason for the incident was the failure of the government to free SLA political prisoners. The bomb itself was estimated to have weighed about 4 pounds.

```
OBJECT event<2>:
    name      IS   "Washington Incident",
    type      IS   "bombing",
    date      IS   "760823",
    time-of-day      IS   "1326",
    location         IS   "Senate",
    locale    IS   "Washington",
    region    IS   "urban",
```

```
host-country    IS   "United States",
announced-claimant IS ("New World Liberation Front"),
when-claimed    IS   ("before-bombing"),
announced-purpose IS ("retaliation against government"),
target  IS   ( "specific individual" ),
target-symbolic-value   IS   "inherently obvious",
target-accessibility    IS   "medium",
victim  IS   ( "person" ),
bomb-weight     IS   "4",
bomb-type       IS   "Molotov-cocktail",
bomb-size       IS   "small",
bomb-delivery   IS   "thrown",
bomb-detonation IS   "impact",
secondary-detonation    IS   "no detonation",
detonation-source       IS   "self-detonating",
explosive-source        IS   ( "home-made" ),
bomb-sophistication     IS   "low",
audacity        IS   "high",
risk-to-perpetrator     IS   "high",
recipient-of-warning    IS   ( "no warning" ),
warning-time    IS   "0",
outcome IS   "malfunction/no detonation",
tactical-objective IS ("publicity/group", "murder"),
strategic-objective     IS   ( "publicity" ),
number-killed   IS   "0",
number-wounded  IS   "0",
primary-damage  IS   "no damage",
secondary-damage        IS   "no damage",
instigator      IS   ("New World Liberation Front"),
perpetrator     IS   ("New World Liberation Front"),
unused  IS   ("disclaimant", "bomb-composition");

OBJECT target<1>:
    name    IS   "Kissinger";
```

Fig. 6 -- A hypothetical bombing incident and the corresponding RITA description

BUILD is an agent for acquiring data from the user. What is also needed is an agent for acquiring judgments or rules from the user. Such an agent would query the user in much the same manner as BUILD and use the information gathered to construct rules or goals about terrorist activities. To present a clearer idea of how such an agent might work a hypothetical user-agent dialogue is shown below.

QUESTION 1: What do you want to know about the event that you do not already know?

ANSWER 1: I want to know the tactical-objective of the event and the political-ideology of the group.

QUESTION 2: What do you think the tactical-objective might be? ("I don't know" is not acceptable.)
ANSWER 2: Extortion.

QUESTION 3: What do you think the political-ideology might be? ("I don't know" is not acceptable.)
ANSWER 3: New left.

QUESTION 4: What pieces of information would you need in order to validate or negate your answers?
ANSWER 4: I would want to know the locale of the event, the target of the event, and the announced-purpose of the event.

QUESTION 5: What is a possible locale for extortion and new left?
ANSWER 5: California.

QUESTION 6: What is a possible target for extortion and new left?
ANSWER 6: Vital systems.

QUESTION 7: What is a possible announced-purpose for extortion and new left?
ANSWER 7: Reduce utility rates for the poor.

RESPONSE: On the basis of your answers, I have constructed the following rule:

RULE 1:
IF: the locale of the event is "California"
& the target of the event is "vital systems"
& the announced-purpose of the event is "reduce utility rates for the poor"
THEN: set the tactical-objective of the event to "extortion-specified philanthropic"
& set the political-ideology of the group to "new left";

QUESTION 8: Does this rule strike you as being irrelevant or inaccurate? If so, rewrite it as one or more rules which you believe are more accurate.

This type of rule acquisition capability would be similar to the one developed for MYCIN (Davis, 1976). The agent and the user would work together to develop and refine the rules, with the

user always having the final say in the matter. Figure 7 shows what some goals acquired this way might look like.

```
GOAL 1:
    IF: the public-opinion OF the current-event
          IS "anti-terrorist"
  THEN: SET the government-response OF the current-event
          TO "harsh crackdown on terrorists";

  [If the public is strongly against the terrorists then the
  government will take harsh measures to suppress the terrorists.]


GOAL 2:
    IF:   the casualty-level OF the current-event IS "high"
            OR THERE IS a campaign WHOSE event-list IS KNOWN
        & the name OF the current-event
            IS IN the event-list OF the campaign
        & the frequency OF the campaign IS "escalating"
            OR the time-between-attacks OF the campaign IS "short"
  THEN: SET the public-opinion OF the current-event
            TO "anti-terrorist";

  [If the current event caused many casualties or is part of
  a campaign whose frequency is escalating or which has a short
  time between attacks then the current event will affect public
  opinion, i.e., the public will become anti-terrorist.]


GOAL 3:
    IF: the number-killed OF the current-event IS greater THAN 0
  THEN: SET the casualty-level OF the current-event TO "high";

  [If 1 or more people are killed the casualty level is
  considered high]


GOAL 4:
    IF: the attacks-per-month OF the campaign IS NOT less THAN 1
  THEN: SET the time-between-attacks OF the campaign TO "short";

  [If the attacks average 1 or more per month they are
  considered frequent.]
```

Fig. 7 -- Goals for deducing information about bombing campaigns

Note that these goals are more complex than the axiomatic-type goals used by BUILD (see Fig. 4). They would be used, in

conjunction with many other similar goals, to deduce information about bombing campaigns.

V. WHAT HAVE WE LEARNED THUS FAR?

The process of formalizing the domain by identifying the important concepts or objects, their attributes, and the associated values is of major importance in the construction of a heuristic model. This is because these basic components are all that are available for describing the complex heuristics or rules that define activity in the domain. If these basic building blocks are too narrow in scope there will be relevant relationships that cannot be expressed. If they are too wide, the system will tend to be inefficient, spending much time processing irrelevant information.

The task of eliciting from analysts judgments that the system can use as "rules" is far more difficult and complex than originally anticipated. "Experts," it appears, have a tendency to state their conclusions and the reasoning behind them in general terms that are too broad for effective machine analysis. It is advantageous to have the machine work at a more basic level, dealing with clearly defined pieces of basic information that it can build into more complex judgments. In contrast, the analyst seldom operates at a basic level. He makes complex judgments rapidly, without laboriously reexamining and restating each step in his reasoning process. The pieces of basic knowledge are assumed and are combined so quickly that it is difficult for him to describe the process. When he examines a problem, he cannot easily articulate each step and may even be unaware of the individual steps taken to reach a solution. He may ascribe to intuition or label a hunch that which is the result of a very

complex reasoning process based upon a large amount of remembered data and experience. In subsequently explaining his conclusion or hunch he will repeat only the major steps often leaving out most of the smaller ones, which may have seemed obvious to him at the time.

Knowing what to consider basic and relevant and not requiring further reevaluation is what makes an analyst an "expert." An economic forecast does not begin with an explanation of why 2 plus 2 equals 4. A political forecast in the United States does not begin with a statement that there are three branches of government and presidential elections every four years. An intelligent discussion about terrorism does not start with the fact that Palestinians are Arabs or that aerial hijacking means taking over an airplane. And when arguing what a particular group might do in a specific situation the analyst seldom states that he is basing his conclusion perhaps in part on the personality of a particular leader, in part on the modus operandi of a particular group, and in part on his recollection of outcomes in similar episodes in the past. This type of information is considered "basic" by the expert. Nor does he discuss why he chose certain past episodes as relevant and discarded others that the non-expert might have chosen, or what specific aspect of the modus operandi is pertinent to the current circumstances. The analyst is not deliberately mysterious about his process of reasoning, nor is his analysis sloppy or incorrect. He simply does not state every single piece of information and every small component part of every judgment he makes. Thus we discovered that the judgments the analysts considered simple and basic were actually complex, often composed of many individual steps that could be elicited only by the

annoying process of repeatedly asking him to justify each statement, including the statements used to clarify previous statements. Obtaining the basic rules the machine needs in order to mimic the reasoning process of the analyst is a difficult and sometimes painful task.

Attempts to extract rules from terrorist experts in the abstract simply by asking them to write all the rules they could think of pertaining to a particular domain did not prove successful for two reasons. First, experts don't usually think of their judgments as being based on a set of "rules," and have trouble putting their ideas into rule form. Second, the rules elicited by this method varied in level of abstraction but generally were all too abstract or complex for the machine to use. They had to be broken down into their component parts, which is something the experts had not been required to do and normally were not accustomed to doing.

We discovered that it was far more useful to elicit rules during or immediately after an actual event in which the analysts were interested and wanted to discuss anyway. The event provided the stimulus for a lively discussion. During the discussion the analysts were asked to offer their opinions or judgments or hunches about some particular aspect of the event. Then they were asked why they felt this to be true. This generally produced a train of rather complex judgments. They were asked to explain how they had made each individual judgment. Each of these produced a train of somewhat less complex judgments which were "pulled apart" by the interrogators, and the process was continued until the critical attributes were identified and basic rules about them articulated.

This by itself was a major achievement. In effect, the extraction process compelled the analyst to examine his own train of thought with an unprecedented degree of rigor. The first and most obvious result was the identification of the attributes of an event or group that were generally agreed upon to be the relevant things to examine. These were the basis for most judgments. The second result was the emergence of some rules about how these attributes interact. But the final result was more than this -- and most important -- the process itself, the extraction of basic rules, sharpened the experts' analytical skills. Regardless of whether these rules could ever be assembled into a system that could in any way approach human reasoning in dealing with complex and ill-defined subjects, being forced to articulate every step along the way to a problematical conclusion was a useful calisthenic for the analysts. It made them more aware, and hence more critical, of their own reasoning; it caused them to examine closely how they arrived at conclusions; and it taught them to look carefully at the spaces between the steps they described. Also, it conditioned them not to overlook things that otherwise might have been ignored, especially in crisis situations where they would be compelled to make snap judgments without time for reflection.

This would seem to indicate a potential utility of this effort that transcends the feasibility of being able to represent the entire domain of terrorism by rules in a minicomputer. It is not necessary for the machine alone to duplicate human thought. Nor is it necessary for the analyst alone to derive from his subject area a set of judgments that can be stored in a computer. It is the interaction of the two that causes an improvement in

the analyst's analytical capabilities, in both ordinary and crisis situations, and results in a tool for the analyst to use.

It is now anticipated that it will require a large number of rules to deal with a small portion of the domain of international terrorism. This problem could be alleviated to some extent by creating many relatively small agents (300 rules or less), each a specialist in some restricted area of the domain. They could all have access to a large agent containing general information about terrorism and would make use of this agent when necessary. Because of the complexity of the task we will need as much computing power as the state of the art can provide. Thus as the power of minicomputers increases, so will the feasibility of modeling complex domains using interactive transaction agents.

VI. APPLICABILITY AND FUTURE DEVELOPMENTS

In addition to the usefulness of heuristic models in improving the skills of analysts working with the problem of international terrorism, it appears that minicomputers with heuristic modeling systems built into them may be extremely useful in the management of crisis situations. Terrorist actions, such as the hijacking of an airliner or the takeover of an embassy, frequently create a crisis situation to which the government must respond. The usual response is to assemble a task force or several task forces at different locations to deal with the crisis. The members of the task force represent different skills (intelligence, public affairs, negotiations, area specialists) and the different agencies and departments that are somehow involved in the episode. The task force decides or offers options to the official who must decide the course of action to take. In dealing with these episodes, particularly hostage situations where lives are at stake, a crisis atmosphere often prevails. Time is critical. (An average barricade-and-hostage incident lasts around 36 hours although some may drag on for days.) Assembling information about the individuals and groups involved is difficult, but this problem can be alleviated with conventional data retrieval systems. Assembling the data that are particularly relevant to a given situation is more difficult. A data retrieval system would not know which special circumstances call for certain actions and which do not. No one has time to read lengthy reports on the group and skilled specialists themselves might not know which reports to look at first. Moreover there is the problem of decisionmaking by

committee. It is difficult to sit around a table for any period of time in silence. People tend to fill the gaps with conversation or discussion about things that may or may not be particularly relevent to the problem at hand. These result in distractions that slow or divert the decisionmaking process.

The system envisioned here first of all provides a straightforward way of proceeding toward a solution. It does this by asking questions about the episode already determined relevant by the experts. In other words, it produces demands for the most relevant information. It summons from its own memory the relevant judgments that have been made about similar groups or events in the past. Furthermore, it enables the committee to explore various options by creating in the machine similar artificial events and changing various responses. The machine can then display likely outcomes to these events. Interaction between the members of the committee and the machine will cause them to examine their own conclusions and decisions in the same manner that it does for the analyst. In sum, the machine imposes a degree of discipline on the course of the discussion, asks the most relevant questions, provides the best judgments of expert analysis, and allows the committee to explore viable options. It also sharpens the committee's skills during the process. It is not essential that the system be entirely accurate or that its performance be flawless. Again, it is the interaction that leads to good sharpened answers, not the machine revealing them from a stored memory.

Remote interactive transaction agents with a capability for examining terrorist incidents would be especially useful in a crisis if they linked together several task forces in a manner

similar to the ARPAnet. This would enable the individuals or groups at different locations to see and participate in each aspect of the decisionmaking. Everyone would see the same video display at the same time. They could intervene with questions or opinions and their interventions could be seen by all of the others.

Such a system also has utility as a storehouse of judgmental information compiled by many expert analysts over a period of time. Indeed, it may be seen as a judgment retrieval system as opposed to a data retrieval system. This would add to the overall amount of expert opinion that is available, would allow interaction between experts at different times and in different locations, and would tend to ameliorate the effect of lost capabilities that result from rotation or unavailability of personnel.

There are many directions in which this reseach can progress in the future. First, we can introduce the notion of certainty factors, probability measures on all the rules and goals that give a subjective estimate of their accuracy. This would provide us with a way of more accurately evaluating the results of our deductions. Second, we can expand the scope of the analysis to other terrorist activities, such as kidnappings or barricade-and-hostage events. The scope can also be widened by including data on more groups and incidents. Third, we can develop agents to help the user recognize and define new rules, as described in Sec. IV. Fourth, we can develop agents that are capable of making predictions about future terrorist activities and use them to test the accuracy of our current formalism. Finally, we can attack the problem of inductive inference,

recognizing regularities in the data and automatically forming rules to describe these regularities.

Appendix A

FORMALIZATION OF A DOMAIN FOR TERRORIST BOMBINGS

event .....

name
[give the event to be defined an identifying name]

type
[type of the terrorist event; e.g. bombing, kidnapping]
- bombing
- barricade-and-hostage
- kidnapping

date
[use the format yymmdd; e.g., 760330]

time-of-day
[a four-digit number; e.g., 0830, 1635]

location
[an object; e.g., embassy, airplane]

locale
[a place; e.g., Paris, Madrid]

region
[environs in which the event occurred]
- urban
- rural
- water
- air

host-country
[name of country in which event occurred]

announced-claimant
[name of person(s) or group claiming credit for the event]

when-claimed
[when were claims announced taking credit for the event?]
- before-bombing
- before-publicity
- before-casualty-report
- after-casualty-report

announced-purpose
[announced reason for carrying out event]

disclaimant
[name of person(s) or group disclaiming credit for the event]

target
[who or what was the immediate object of the attack?]
   specific individual
   representative individual
   random individual
   political property
   corporate headquarters or offices
   other commercial property
   private property
   vital systems
   head of state

target-symbolic-value
[what symbolic value did the object of the attack have?]
   inherently obvious
   specified by perpetrator
   none apparent

target-accessibility
[how accessible was the object of the attact to the perpetrators]
   low
   medium
   high

victim
[who or what was the event intended to affect?]
   person
   organization
   government
   ideal

bomb-weight
[estimated weight in pounds; e.g., 10]

bomb-type
[which type of bomb was used in the event?]
   Molotov-cocktail
   other incendiary
   letter or parcel
   pipe
   small explosive charge
   shopping bag
   grenade
   car
   land mine

bomb-size
[estimate the size of the bomb: small, medium, or large]
   small
   medium
   large

bomb-delivery
[how did the bomb arrive at its destination?]
   planted
   delivered
   thrown

bomb-detonation
[what method was used to detonate the bomb?]

command-line
command-radio
impact
time
mechanical
electrical
chemical
fuse
x-ray
photoelectric
altitude

secondary-detonation
[what type of backup detonation device did the bomb have?]
no detonation
command-line
command-radio
impact
time
mechanical
electrical
chemical
x-ray
photoelectric
altitude

detonation-source
[who detonated the bomb?]
victim
perpetrator
self-detonating

bomb-composition
[what type of explosive material was used in the bomb?]

explosive-source
[source of the explosive material: commercial, military, etc.]
commercial
military
home-made

bomb-sophistication
[give a subjective estimate of the sophistication of the bomb]
low
medium
high

audacity
[give an appraisal of how brazen and audacious the incident was]
low
medium
high

risk-to-perpetrator
[assess the risk to the perpetrators of capture or death]
low
medium
high

recipient-of-warning

[who received prior warning of the bomb?]
   media
   target
   police
   no warning

warning-time
[minutes before the expected detonation the warning was received]

outcome
[what happened to the bomb?]
   detonated
   discovered/defused
   malfunction/no detonation

tactical-objective
[what was the immediate goal of the bombing?]
   murder
   punishment/retaliation
   property damage-no casualties
   property damage-casualties
   extortion-payoff to perpetrator
   extortion-unspecified philanthropic
   extortion-specified philanthropic
   disruption
   publicity/group
   publicity/commemorative

strategic-objective
[what was the more global objective of the instigator?]
   increase-pain-level
   halt-operation
   publicity
   discredit-government
   political-victory
   military-victory
   defense of the group/movement
   intergroup rivalry

number-killed
[number of people killed during incident]

number-wounded
[number of people wounded during the incident]

primary-damage
[what was the damage resulting directly from the bomb itself]
   no damage
   low
   medium
   high

secondary-damage
[what was the damage caused in the aftermath of the bomb]
   no damage
   low
   medium
   high

instigator

[name of person(s) or group who planned the event]

perpetrator
[name of person(s) or group carrying out the event]

group .....

name
[type the full proper name of the group and the acronym, if any]

political-ideology
[political persuasion or connections of group]
fascist
right-wing
Moscow-line communist
independent communist
Trotskyite
Maoist
anarchist
new left

nationality
[country or origin or primary base]

age
[how many years has the group been in existence?]

objective
[what are the strategic objectives of the group?]
enlarge the group
expansion
overthrow of government
discredit the system
territorial independence
limited reforms or concessions
vigilante-like defense

action-frequency
[the approximate number of incidents perpetrated per year]

past-incidents
[type the names of the past incidents instigated by the group]

affiliations
[names of allied groups]

stage-of-development
[stages 1 to 5 allowed]

ethnic-composition
[major nationality or ethnic composition represented in the group]

educational-composition
[what kind of education do the group members have?]
uneducated
high school

university

supporters
[groups, sectors, and organizations that aid and abet the group.]
other terrorist groups.
university-educated population sectors
working-class sectors
minority groups
foreign governments
nobody

level-of-sophistication
[sophistication of the group in terms of planning and execution]
low
medium
high

tactics
[the typical types of events perpetrated by the group]
bombing
barricade-hostage
kidnapping
assassinations
highjackings
armed assaults
robbery
dissemination of propaganda
sabotage

type
[characterize the goal of the group in terms of the given categories]
separatist/irredentist
national revolutionary
international revolutionary

activity-span
[how many years has the group been active at terrorism?]

targets
[the types of objects or persons typically attacked]
specific individual
representative individual
random individual
government property
political property
corporate headquarters or offices
other commercial property
private property
vital systems

members
[type names of members: first name (alias) last name]

headquarter-location
[where the decisonmakers are located, e.g. Berkeley, Beirut]

territory-of-operations
[the territorial area over which the group spreads its operations]
international
national

regional
state
local

level-of-support
[extent of support for the group: e.g., local, regional]
insignificant
local
regional
national
international

type-of-support
[expressions of support for the group]
passive sympathy
limited help
total commitment

person .....

name
[name of the person]

nationality
[name of country of national origin]

sex
[male or female]

age
[approximate age in years]

occupation
[what does the person do for a living?]
police/military
government
corporate official
political-figure
private-citizen

group-membership
[organizational affiliation]

terrorism-role
[how is this person connected with terrorist activities?]
target
perpetrator
instigator
hostage

Appendix B

PROTOCOL OF AN EXPERT USING THE BUILD AGENT

```
% rita use.build.goal

use.build.goal:

   Fill-in-the-blanks questions can be answered with either
a single reply such as "Fred Jones", or with a list of
replies each separated by a comma followed by a space, i.e.,
"Fred Jones, Mary Smith, John Doe".  To indicate that
you don't know just type <carriage return>.

   For questions of the form: "The <a> of the <b> is...",
typing "unknown" is equivalent to typing <carriage return>.
For questions of the form: "Other <a>s of the <b>...", typing
"none" is equivalent to typing <carriage return>.

   Yes-or-no questions can be answered with either
"y", or "yes" for yes; "n", "no", or carriage
return for no; or a single answer, such as
"school building" for a qualified yes.

   You may type "?" in response to any question to
elicit information about the type of answer required.

   To terminate the session type "nothing" in response to
the "What is to be defined..." question.  Then to see the
data base that was just created type "display all objects;".

The most current event took place in Washington D.C.
at about 1:26 pm in the Senate Office Building on August
23, 1976.  As Secretary of State Henry Kissinger was
exiting the front door a lone individual hurled a
Molotov-cocktail at him and then fled on foot.  The bomb
failed to explode and no one was hurt.  Two minutes before
the attempted bombing a woman who identified herself as
a member of the New World Liberation Front called the
Washington Post and claimed credit for the act.  She said
the reason for the incident was the failure of the
government to free SLA political prisoners.  The bomb
itself was estimated to have weighed about 4 pounds.


What is to be defined...event

The name of the event is: Washington Incident

Is the type of the event :
  bombing? yes

The date of the event is: ?

[use the format yymmdd; e.g., 760330]

The date of the event is: 760823
```

The time-of-day of the event is: ?

[a four-digit number; e.g., 0830, 1635]

The time-of-day of the event is: 1326

The location of the event is: Senate

The locale of the event is: Washington

Is the region of the event :
  urban? yes

The host-country of the event is: United States

The announced-claimant of the event is: ?

[name of person(s) or group claiming credit
for the event]

The announced-claimant of the event is: New World Liberation Front

Is the when-claimed of the event :
  before-bombing? ?

[when were claims announced taking credit for the incident? Give one answer for each announced claim]

  before-bombing? yes
  before-publicity? no
  before-casualty-report? no
  after-casualty-report? no

Other when-claimed(s) of the event...

The announced-purpose of the event is: retaliation against government

The disclaimant of the event is:

Is the target of the event :
  specific individual? Kissinger
  representative individual? no
  random individual? no
  political property? no
  corporate headquarters or offices? no
  other commercial property? no
  private property? no
  vital systems? no
  head of state? no

Other target(s) of the event...

Is the target-symbolic-value of the event :
  inherently obvious? yes

Is the target-accessibility of the event :
  low? no
  medium? yes

```
[I deduce that the victim(s) of the event includes:
 ( "person")]

The bomb-weight of the event is: ?

[estimated weight in pounds; e.g., 10]

The bomb-weight of the event is: 4

Is the bomb-type of the event :
  Molotov-cocktail? yes
[I deduce that the bomb-size of the event is "small"]
[I deduce that the bomb-delivery of the event is "thrown"]
[I deduce that the bomb-detonation of the event
 is "impact"]
[I deduce that the secondary-detonation of the event
 is "no detonation"]
[I deduce that the detonation-source of the event
 is "self-detonating"]

The bomb-composition of the event is:
[I deduce that the explosive-source(s) of the event includes:
 ( "home-made")]
[I deduce that the bomb-sophistication of the event is "low"]

Is the audacity of the event :
  low? no
  medium? no
  high? yes

Is the risk-to-perpetrator of the event :
  low? no
  medium? no
  high? yes

Is the recipient-of-warning of the event :
  media? no
  target? no
  police? no
  no warning? yes

Other recipient-of-warning(s) of the event...
[I deduce that the warning-time of the event is "0"]

Is the outcome of the event :
  detonated? no
  discovered/defused? no
  malfunction/no detonation? yes
[I deduce that the tactical-objective(s) of the event includes:
 ( "publicity/group",
                "murder")]
[I deduce that the strategic-objective(s) of the event includes:
 ( "publicity")]
[I deduce that the number-killed of the event is "0"]
[I deduce that the number-wounded of the event is "0"]
[I deduce that the primary-damage of the event is "no damage"]
[I deduce that the secondary-damage of the event is "no damage"]
[I deduce that the instigator(s) of the event includes:
 ( "New World Liberation Front")]
[I deduce that the perpetrator(s) of the event includes:
```

```
 ( "New World Liberation Front")]

Finished with the event


What is to be defined...nothing

The following data were added to
file bomb.data in your file area:


[OBJECTS:]

OBJECT event<2>:
        name    IS   "Washington Incident",
        type    IS   "bombing",
        date    IS   "760823",
        time-of-day     IS   "1326",
        location        IS   "Senate",
        locale  IS   "Washington",
        region  IS   "urban",
        host-country    IS   "United States",
        announced-claimant IS ("New World Liberation Front"),
        when-claimed    IS   ( "before-bombing" ),
        announced-purpose IS ("retaliation against government"),
        target  IS   ( "specific individual" ),
        target-symbolic-value   IS   "inherently obvious",
        target-accessibility    IS   "medium",
        victim  IS   ( "person" ),
        bomb-weight     IS   "4",
        bomb-type       IS   "Molotov-cocktail",
        bomb-size       IS   "small",
        bomb-delivery   IS   "thrown",
        bomb-detonation IS   "impact",
        secondary-detonation    IS   "no detonation",
        detonation-source       IS   "self-detonating",
        explosive-source        IS   ( "home-made" ),
        bomb-sophistication     IS   "low",
        audacity        IS   "high",
        risk-to-perpetrator     IS   "high",
        recipient-of-warning    IS   ( "no warning" ),
        warning-time    IS   "0",
        outcome IS   "malfunction/no detonation",
        tactical-objective      IS   ("publicity/group", "murder"),
        strategic-objective     IS   ("publicity"),
        number-killed   IS   "0",
        number-wounded  IS   "0",
        primary-damage  IS   "no damage",
        secondary-damage        IS   "no damage",
        instigator      IS   ( "New World Liberation Front" ),
        perpetrator     IS   ( "New World Liberation Front" ),
        unused  IS   ( "disclaimant", "bomb-composition" );


OBJECT target<1>:
        name    IS   "Kissinger";



exiting.
%
```

BIBLIOGRAPHY

Anderson, R. H., and J. J. Gillogy, Rand Intelligent Terminal Agent (RITA): Design Philosophy, The Rand Corporation, R-1809-ARPA, February 1976.

Anderson, R. H., and J. J. Gillogly, Rand Intelligent Terminal Agent (RITA): Reference Manual, The Rand Corporation, R-1808-ARPA, 1976 (forthcoming).

Davis, Randall, Applications of Meta Level Knowledge to the Contruction, Maintenance, and Use of Large Knowledge Bases, Stanford University, Artificial Intelligence Laboratory, Memo AIM-283, 1976.

Davis, Randall, et al., Production Rules as a Representation for a Knowledge-Based Consultation Program, Stanford University, Artificial Intelligence Laboratory, Memo AIM-266, 1975.

Feigenbaum, E. A., et al., On Generality and Problem Solving: A Case Study Using the DENDRAL Program, in B. Meltzer and D. Michie (eds.), Machine Intelligence 6, Edinburgh University Press. (Also in Stanford University, Artificial Intelligence Laboratory, Memo AIM-131, August 1970.)

Lenat, Douglas, B., AM: An Artificial Intelligence Approach to Discovery in Mathematics as Heuristic Search, Stanford University, Artificial Intelligence Laboratory, Memo AIM-286, Ph.D. dissertation, July, 1976.

Newell, A., and H. A. Simon, Human Problem Solving, Prentice-Hall, 1972.

Newell, A., A Theoretical Exploration of Mechanisms for Coding the Stimulus, in A. W. Melton and E. Martin (eds.), Coding Processes in Human Memory, Winston and Sons, Washington D.C., 1972.

Newell, A., Production Systems: Models of Control Structures, in W. C. Chase (ed.), Visual Information Processing, 1973.

Shortliffe, E. H., MYCIN: A Rule Based Computer Program for Advising Physicians Regarding Antimicrobial Therapy Selection, Stanford University, Ph.D. dissertation, 1974.

Shortliffe, E. H., and B. G. Buchanan, A Model of Inexact Reasoning in Medicine, Mathematical Biosciences, 23, 1975, pp. 351-379.

Waterman, D. A., Generalization Learning Techniques for Automating the Learning of Heuristics, Artificial Intelligence, 1, 1970, pp. 121-170.

Waterman, D. A., and A. Newell, Preliminary Results with a System for Automatic Protocol Analysis, Carnegie-Mellon University, Computer Science Department, 1973.

Waterman, D. A., Adaptive Production Systems, International Joint Conference on Artificial Intelligence-4 Proceedings, Vol. 1, September 1975, pp. 296-303.

Waterman, D. A., and A. Newell, PAS-II: An Interactive Task-Free Version of an Automatic Protocol Analysis System, IEEE Transactions, April 1976.

Waterman, D. A., Serial Pattern Acquisition: A Production System Approach, Proceedings of the 1976 Workshop on Pattern Recognition and Artificial Intelligence, August 1976.

Waterman, D. A., Rule-Directed Interactive Transaction Agents: An Approach to Knowledge Acquisition, The Rand Corporation, draft report, R-2171, (forthcoming).