

Technical Report TR-1708

September 1986

## An Evaluation of Expert Systems for Software Engineering Management \*

Connie Loggia Ramsey and Victor R. Basili

Department of Computer Science  
University of Maryland at College Park

### ABSTRACT

Although the field of software engineering is relatively new, it can benefit from the use of expert systems. Four prototype expert systems have been developed to aid in software engineering management. Given the values for certain metrics, these systems will provide interpretations which explain any abnormal patterns of these values during the development of a software project. The four expert systems, which solve the same problem, were built using two different approaches to knowledge acquisition, a bottom-up approach and a top-down approach, and two different expert system methods, rule-based deduction and frame-based abduction. A comparison was performed to see which methods better suit the needs of this field. It was found that the bottom-up approach led to better results than did the top-down approach, and the rule-based deduction systems using simple rules provided more complete and correct solutions than did the frame-based abduction systems.

---

\* Research supported in part by the National Aeronautics and Space Administration Grant NSG-5123 to the University of Maryland. Computer support provided in part by the Computer Science Center of the University of Maryland.

THE UNIVERSITY OF CHICAGO

DEPARTMENT OF CHEMISTRY

PHYSICAL CHEMISTRY

PROFESSOR [Name]

REPORT

On the [Title of Report]

[Date]

Submitted by [Name]

## 1. INTRODUCTION

The importance of expert systems is growing in industrial, medical, scientific, and other fields. Several major reasons for this are: (1) the necessity of handling an overwhelming amount of knowledge in these areas, (2) the potential of expert systems to train new experts, (3) the potential to learn more about a field while organizing knowledge for the development of expert systems, (4) cost reductions sometimes provided by expert systems, and (5) the desire to capture corporate knowledge so it is not lost as personnel changes.

Although the field of software engineering is still relatively new, it can certainly benefit from the use of expert systems because of the ability to learn from them. The development of an expert system for software engineering provides a learning experience by forcing the knowledge engineer to develop and organize relationships between various pieces of knowledge, such as metrics and their interpretations, and also by forcing us to realize how much we do not know yet.

On another level, the expert systems in this field can be used to train and help people, including software managers. They can contain general software engineering knowledge as well as a history of information from a particular software development environment which can be particularly helpful to inexperienced managers and developers.

Since software engineering is still such a new field, special attention has to be given to basic research issues concerning the development of expert systems in this field. The high level goals of this project were to determine: (1) Are expert systems for software engineering feasible at this time? (2) What methodology should be used for knowledge acquisition? (3) What type of expert system methodology best suits software engineering? (4) Do the experts themselves agree on the information to be used? (5) Are certain software environments more suited for expert systems than others? (6) Are we ready to develop systems with environment-independent, general truths? (7) What information should be included in the systems?

This paper will focus on several prototype expert systems, collectively named ARROWSMITH-P.\* Earlier versions of these expert systems are described in [Basili85]. This represents an early attempt at defining the process of creating expert systems for software engineering. ARROWSMITH-P is intended to aid the manager of a software development project in an automated manner. The goal of these systems is to help detect and assess problems as early as possible. The systems work as follows. First, it is determined whether or not a software project is following normal development patterns by comparing measures such as programmer hours per line of source code against historical, environment-specific baselines of such measures. Then, the "manifestations" detected by this comparison, such as an abnormally high rate of programmer hours per line of source code, serve as input to each expert system, and each system attempts to determine the reasons, such as *low productivity*, for any abnormal software development patterns. These systems can be updated as the environment changes and as more is learned in the field of software engineering.

---

\* Martin Arrowsmith, created by Sinclair Lewis in the novel *Arrowsmith*, was in constant search of truth in scientific fields. The "P" stands for Prototype.

The rest of this paper is organized as follows. Section 2 provides a brief overview of the underlying methodology used to build the expert systems discussed in this paper. Section 3 details the implementations of ARROWSMITH-P, and Section 4 discusses the issues and problems associated with this process. Section 5 furnishes the details for the evaluation of the expert systems. Section 6 then discusses results and conclusions from the development and testing of the expert systems. Finally, Section 7 discusses current and future research needs.

## 2. BACKGROUND

In general, an expert system consists of two basic components, a domain-specific knowledge base and a domain-independent inference mechanism. The knowledge base consists of data structures which represent general problem-solving information for some application area. The inference mechanism uses the information in the knowledge base along with problem-specific input data to generate useful information about a specific case.

The set of expert systems in ARROWSMITH-P was constructed using KMS [Reggia82a], an experimental domain-independent expert system generator which can be used to build rule-based, frame-based and Bayesian systems. The ARROWSMITH-P systems were built using two different methods: rule-based deduction and frame-based abduction. These two methods are briefly described below.

### 2.1. Rule-Based Deduction

A common method for expert systems, and essentially the "standard" in AI today, is rule-based deduction. In this approach, domain-specific problem-solving knowledge is represented in rules which are basically of the form:

"IF <antecedents> THEN <consequents>",

although the exact syntax used may be quite different (e.g., PROLOG). If the antecedents of such a rule are determined to be true, then it logically follows that the consequents are also true. Note that these rules are not branching points in a program, but are non-procedural statements of fact.

The inference mechanism consists of a rule interpreter which, when given a specific set of problem features, determines applicable rules and applies them in some specified order to reach conclusions about the case at hand. Rule-based deduction can be performed in a variety of ways, and rules can be chained together to make multiple-step deductions. (For a fuller description, see [Hayes-Roth78].) In addition, in many systems one can attach "certainty factors" to rules to capture probabilistic information, and a variety of mechanisms can be used to propagate certainty measures during problem solving. MYCIN [Shortliffe76] and PROSPECTOR [Campbell82] are two well-known examples of expert systems which incorporate rule-based deduction.

### 2.2. Frame-Based Abduction

Another important method for implementing expert systems is frame-based abduction. Here, the domain-specific problem-solving knowledge is represented in descriptive

“frames” of information [Minsky75], and inference is typically based on hypothesize-and-test cycles which model human reasoning as follows. Given one or more initial problem features, the expert system generates a set of potential hypotheses or “causes” which can explain the problem features. These hypotheses are then tested by (1) the use of various procedures which measure their ability to account for the known features, and (2) the generation of new questions which will help to discriminate among the most likely hypotheses. This cycle is then repeated with the additional information acquired. This type of reasoning is used in diagnostic problem solving (see [Reggia82b] for a review). INTERNIST [Miller82], KMS.HT [Reggia82a], [Reggia83a], PIP [Pauker76], and IDT [Shubin82] are typical systems using frame-based abduction.

In order to simulate hypothesize-and-test reasoning, KMS employs a generalized set covering model in which there is a universe of all possible manifestations (symptoms) and a universe which contains all possible causes (disorders). For each possible cause, there is a set of manifestations which that cause can explain. Likewise, for each possible manifestation, there is a set of causes which could explain the manifestation. Given a diagnostic problem with a specific set of manifestations which are present, the inference mechanism finds all sets of causes with minimum cardinality\*\* which could explain (cover) all of the manifestations. For a more detailed explanation of the theory underlying this approach and the problem-solving algorithms, see [Reggia83a], [Reggia83b], [Nau84], [Peng84].

### 3. IMPLEMENTATIONS

In this section, we will first present the methodology developed for building expert systems for software engineering. Then we will discuss the actual implementations of ARROWSMITH-P.

#### 3.1. Methodology

The following two methodologies of knowledge acquisition for constructing expert systems for software engineering management were developed. They can best be described as a bottom-up methodology and a top-down methodology. (An earlier version of the bottom-up reasoning was presented in [Doerflinger83].)

##### 3.1.1. Bottom-Up Methodology

Given a homogeneous environment, it is possible to produce historical, environment-specific baselines of normalized metrics from the data of past software projects. Normalized metrics are derived by comparing variables such as programmer hours and lines of code against each other. This is done so that influences such as the size of the individual project are factored out. The baseline for each metric is defined as the average value of that metric for the past projects at various discrete time intervals (such

---

\*\* Ockham's razor, which states that the simplest explanation is usually the correct one, together with the assumption of independence among causes motivate the requirement of minimum cardinality.

as early coding or acceptance testing). Only those metrics which exhibit baselines with reasonable standard deviations should be used; too little variety in the values of the measures proves uninteresting, while too much variety is not very meaningful. In addition, one ideally wants a relatively small number of meaningful metrics whose values are easily obtainable.

Next, experts can determine interpretations, such as *unstable specifications* or *good testing*, which would explain any significant deviation (more than one standard deviation less than or greater than the average) of a particular metric from the historical baseline. The deviation of some metric can be thought of as a manifestation or symptom which can be "diagnosed" as certain interpretations or causes. Furthermore, these relationships between interpretations and manifestations should be made time-line specific because, for example, an interpretation during early coding might not be valid during acceptance testing. In addition, measures to indicate how certain one is that the deviation of a particular metric has resulted from a particular interpretation can be included.

The approach, described above, can be classified as a bottom-up approach because it seems to go in the opposite direction of cause-and-effect. First the symptoms (deviant metric values) that something is abnormal are explored, and then the underlying interpretations or diagnoses of the abnormalities are developed. This approach to knowledge acquisition is reasonable in a homogeneous environment because the metrics are homogeneous, and deviations are indicative that something is wrong. However, this approach contrasts with the development of expert systems in other fields, such as medicine, which typically use a top-down approach.

### **3.1.2. Top-Down Methodology**

A top-down approach to knowledge acquisition can be similar to the bottom-up approach in that the same manifestations and causes can be used. However, it would first define the various interpretations or diagnoses and then associate the metrics which would be likely to have abnormal values for each interpretation.

Using the top-down approach, the experts view the knowledge from a different perspective when defining the relationships that exist between the interpretations and manifestations. This approach can be seen as a more general approach than the bottom-up approach is to knowledge acquisition in the field of software engineering. In the bottom-up methodology, the metrics are analyzed first and these are, by their nature, environment-specific. The focus is automatically limited to the specific environment. Conversely, in the top-down methodology, the experts think first of the causes or interpretations and then indicate the effects or likely metrics which would show deviant values if a certain interpretation existed. This generalizes the problem across environments somewhat because the emphasis seems to be switched to the interpretations which can be universal.

### **3.1.3. Using the Expert Systems**

Once the expert systems have been developed, the input to each expert system would then consist of those metrics from a current project which deviate from a historical baseline of the same metrics at the same time of development for similar projects. The knowledge bases consist of information about various potential causes, such as *poor testing* or *unstable specifications*, for any abnormally high or low measures, and the expert system provides explanations for any abnormal software development patterns.

### 3.2. Actual Implementations

ARROWSMITH-P consists of four separate expert systems, one using a bottom-up approach to knowledge acquisition and rule-based deduction, a second using the bottom-up approach and frame-based abduction, a third using a top-down approach to knowledge acquisition and rule-based deduction, and a fourth using the top-down approach and frame-based abduction.

The bottom-up methodology described above was based on previous research conducted on the NASA/Goddard Space Flight Center Software Engineering Laboratory (SEL) environment [Doerflinger83]. Since the SEL environment is homogeneous, it was possible to produce historical, environment-specific baselines of normalized metrics from the highly reliable data of nine software projects. (See [Basili77], [Basili84b], [Basili78], [Card82], [SEL82] for fuller descriptions of the SEL environment.)

The bottom-up development was performed first, and nine metrics, derived from five variables, proved satisfactory, exhibiting baselines with reasonable standard deviations. The metrics are displayed in Table 1. These same metrics were later used during the top-down development to ensure consistency and to allow a comparative study to be performed. The time-line for the baselines was divided (after a slight modification) into the following five discrete intervals: early code, middle code, late code, systems test, and acceptance test.

The initial sets of interpretations and the relationships between the interpretations and the abnormal values of metrics were mainly derived from two experts who have had a great deal of experience in this field and particularly in the SEL environment [Doerflinger83]. During the bottom-up development of ARROWSMITH-P, mainly one of these experts modified the existing sets and made them time-line specific. In addition, measures to indicate how certain one is that the interpretation and the abnormal metric value are connected were included. During the top-down development, the same two experts were again asked to provide the relationships for all five time phases, and the intersection of their responses was used for the expert systems. Some of their other indicated relationships were used as well; when the experts did not agree on a relationship, we discussed the situation to understand the reasoning behind the relationship and to see how certain an expert felt about the relationship. The list of interpretations used and tested in the bottom-up and top-down expert systems is displayed in Table 2.

TABLE 1 - METRICS USED IN EXPERT SYSTEM

- Computer Runs per Line of Source Code
- Computer Time per Line of Source Code
- Software Changes per Line of Source Code
- Programmer Hours per Line of Source Code
- Computer Time per Computer Run
- Software Changes per Computer Run
- Programmer Hours per Computer Run
- Computer Time per Software Change
- Programmer Hours per Software Change

(Other interpretations were used as well, but these could not be tested. See [Basili85] for the complete list.)

As stated previously, two different expert system methods were used to build the expert systems for this application in order to determine which method better suits the needs of this field. The two methods used were rule-based deduction and frame-based abduction which were described in Section 2. In the rule-based systems, the rules are of the form "IF manifestations THEN interpretations," while in the frame-based systems, there is one frame (containing a list of manifestations) for each interpretation. Please note that these formats are independent of whether the relationships between manifestations and interpretations were defined using a bottom-up or a top-down approach to knowledge acquisition. The rule-based and frame-based systems which used the bottom-up approach were intentionally built to be as consistent with one another as possible. The causes and manifestations used were identical in both cases, as were the relationships between them. The same was true for the two expert systems which employed the top-down approach. However, the certainty factors attached to the rules and the measures of likelihood in the frames could not be directly translated to each other so some of these measures were omitted. For example, within the bottom-up approach we were relatively certain that an abnormally high value of computer time per software change is caused by *good, reliable code* so this was given a certainty factor of 0.75.

TABLE 2 - INTERPRETATIONS USED IN EXPERT SYSTEM

Unstable Specifications
Low Productivity
High Productivity
High Complexity or Tough Problem
High Complexity or Compute Bound Algorithms Run or Tested
Low Complexity
Simple System
Error Prone Code
Good Solid and Reliable Code
Large Portion of Reused Code
Lots of Testing
Little Testing
Good Testing or Good Test Plan
Lack of Thorough Testing
Poor Testing Program
Changes Hard to Make
Loose Configuration Management or Unstructured Development
Tight Configuration Management or Control
Computer Problems or Inaccessibility or Environmental Constraints
Lots of Terminal Jockeys

Note - In the top-down systems, *low complexity* and *simple system* were combined into one interpretation as were *lack of thorough testing* and *poor testing program*. In addition, the *tight management* interpretation was removed.



However, if that particular metric appears abnormally high very infrequently and that particular interpretation is common, then we would not be able to state that *good, reliable code* generally results in an abnormally high value of computer time per software change. (For a discussion of similar problems see [Ramsey86].) Figure 1 shows a sample section of a rule-based and a frame-based knowledge base. Example sessions with the expert systems are provided in Appendix 1.

#### 4. RESEARCH ISSUES AND PROBLEMS

The field of expert systems is relatively new, and therefore, the development process of expert systems still faces many problems. The selection of which method to use for building them is not generally clear, although an attempt has been made to provide guidelines for the selection of an appropriate method in [Ramsey86]. Furthermore, most expert systems are shallow in nature and cannot handle temporal or spatial information well.

In addition to general problems, negative effects are compounded when the knowledge to be included in such systems is incomplete. The science of software engineering is not well-defined yet, and therefore many details about the relationships between various components is often unclear. The experts themselves may not even agree on the information used in the expert systems. As a result, the knowledge base of any expert system developed in this field is particularly exploratory and prototypical in nature. This is in contrast to expert systems developed in established fields such as medicine where the information contained in the knowledge base is based on many years of experience.

Due to the uncertainty of the data in the knowledge base for a field such as software engineering, one must deal with the issues of completeness versus correctness and completeness versus minimality. When dealing with a diagnostic problem, the more certain one is of relationships between causes and manifestations, the more exact the answer can be, ultimately leading to the one correct answer. However, when dealing with very uncertain relationships, it is preferable to list many outcomes so as to avoid missing the correct explanation, and to let the experienced person using the expert system decide what the correct explanation really is. Therefore, rules with simple antecedents were used in the rule-based deduction systems (see Figure 1a) because the more involved patterns needed for complex antecedents are not yet known. If one tried to "guess" what these patterns are without actually being certain, this would lead to incomplete solutions which miss some of the correct interpretations. For example, a high value for computer runs per line of code, a high value for computer time per line of code, and a high value for programmer hours per line of code are all indications of *low productivity*. So, we might construct the following rule for this pattern:

IF Computer Runs per Line of Code is above normal, and Computer Time per Line of Code is above normal, and Programmer Hours per Line of Code is above normal  
THEN the interpretation is *Low Productivity*.

However, what if it turns out that computer time per line of code is almost never above normal? Then this rule will almost never succeed, and we will miss the interpretation of *low productivity* even if it happens to be true.

ATTRIBUTES:

```
/* INPUT ATTRIBUTES */  
COMPUTER RUNS PER LINE OF SOURCE CODE (SGL):  
  ABOVE NORMAL,  
  NORMAL,  
  BELOW NORMAL.
```

```
/* INFERRED ATTRIBUTE */  
INTERPRETATION (MLT):  
  UNSTABLE SPECIFICATIONS  
  LOW PRODUCTIVITY  
  HIGH PRODUCTIVITY  
  GOOD TESTING OR GOOD TEST PLAN
```

RULES:

```
CRLC1 IF COMPUTER RUNS PER LINE OF CODE = ABOVE NORMAL,  
  & TIME = EARLY CODING  
  THEN INTERPRETATION = LOW PRODUCTIVITY <0.25>,  
  & INTERPRETATION = ERROR PRONE CODE <0.75>.
```

```
SCLC3 IF SOFTWARE CHANGES PER LINE OF CODE = ABOVE NORMAL,  
  & TIME = LATE CODING  
  THEN INTERPRETATION = GOOD TESTING OR GOOD TEST PLAN <0.25>,  
  & INTERPRETATION = ERROR PRONE CODE <0.75>.
```

Figure 1a - Small Section of Rule-Based Deduction Expert System.

ATTRIBUTES:

```
/* INPUT ATTRIBUTES */  
COMPUTER RUNS PER LINE OF SOURCE CODE (SGL):  
  ABOVE NORMAL,  
  * NORMAL,  
  BELOW NORMAL.
```

```
/* INFERRED ATTRIBUTE - FRAMES */  
INTERPRETATION (MLT):  
  LOW PRODUCTIVITY  
  [DESCRIPTION:  
    COMPUTER RUNS PER LINE OF CODE = ABOVE NORMAL;  
    COMPUTER TIME PER LINE OF CODE = ABOVE NORMAL;  
    PROGRAMMER HOURS PER LINE OF CODE = ABOVE NORMAL ],  
  GOOD TESTING OR GOOD TEST PLAN  
  [DESCRIPTION:  
    SOFTWARE CHANGES PER LINE OF CODE = ABOVE NORMAL;  
    SOFTWARE CHANGES PER COMPUTER RUN = ABOVE NORMAL;  
    COMPUTER TIME PER SOFTWARE CHANGE = BELOW NORMAL;  
    PROGRAMMER HOURS PER SOFTWARE CHANGE = BELOW NORMAL ],
```

Figure 1b - Small Section of Frame-Based Abduction Expert System

This issue also leads to concern in the frame-based abduction systems which provide all answers of minimum cardinality. The inference mechanism works very well for most diagnostic problem solving, but one must be cautiously aware of the fact that not all possible explanations are provided by this expert system. For example, if an abnormally high value of computer runs per line of code and an abnormally low value of programmer hours per software change can be explained by the combination of two interpretations, *low productivity* and *good testing*, and also by a single interpretation, *error prone code* alone, then only the single interpretation will be provided by this system. This is because the single interpretation has a lower cardinality than the two interpretations together.

One final, but very important, fact should be noted here. ARROWSMITH-P was built using the data from one particular homogeneous environment. Therefore, the information in the knowledge base reflects this one environment and would not be transportable to other environments. However, the ideas and methods used to build ARROWSMITH-P are transportable, and that is what is important.

## 5. EVALUATION OF EXPERT SYSTEMS

### 5.1. Methods of Evaluation

ARROWSMITH-P has been evaluated in several ways. The correctness of the systems was measured by comparing the interpretations provided by the expert systems against what actually happened during the development of the projects, thereby obtaining a measure of agreement. This analysis was performed for ten projects (the original nine plus a newer project which was completed after the development of the expert systems) in all five time phases for each of the four expert systems. Each of the original nine projects was compared against historical baselines of the remaining eight projects to determine abnormal metric values, and the tenth project, which was tested later, was compared against the original nine. A total set of 50 cases was tested on each of the four expert systems.

The actual results of what took place during development were gathered from information in another section of the database, mostly from subjective evaluation forms and project statistics forms. The subjective evaluation form contains mostly subjective information (such as a rating of the programming team's performance) and some objective numbers (such as total number of errors) concerning the project's overall development. Since the vast majority of the ratings in the subjective evaluation form is not divided by phase of the project, there probably exist some discrepancies between the results indicated in the forms and the actual interpretations for a particular phase. However, these are the closest data that are available, so we must assume that most of the interpretations for each phase are similar to the interpretations for the entire project. In addition, some of the interpretations derived from analyzing the data in the database were very evident, while others were somewhat uncertain. Therefore, these two classes were partitioned in the analysis of agreement between the expert systems and the information in the database.

The results from the expert systems were also analyzed statistically by using a Kappa statistic test [Spitzer67] [Cohen68] on each interpretation. The Kappa statistic

determines whether the results are better or worse than chance agreement. It takes into account the number of correct answers and the number of incorrect answers with respect to each interpretation, and it determines the amount of agreement which can be attributable to chance alone. The formula for the Kappa statistic is:

$$K = \frac{P_o - P_c}{1 - P_c}$$

where  $P_o$  is the observed proportion of agreement, and  $P_c$  is the proportion of agreement expected by chance. A value of 1 for  $K$  indicates perfect agreement, a value of 0 indicates that the results can be due to chance alone, and a value less than 0 indicates worse than chance agreement. The Kappa statistic was used for each interpretation in each of the four expert systems. This was done to determine whether certain interpretations are better understood than others.

In addition to testing the performance of the expert systems, an analysis was performed to compare the information provided by the two experts for the systems. This was performed by comparing the relationships indicated by each of the experts against each other and also by comparing the relationships indicated in the bottom-up systems against those indicated using the top-down approach.

## 5.2. Results

The first results we would like to discuss are those comparing information provided by the experts. This is essential because the expert systems can only perform as well as the knowledge contained in the systems permits. The experts were asked to fill in grids (one for each time phase for the bottom-up approach and one for each time phase for the top-down approach) indicating the relationships between the interpretations and the manifestations as described in Section 3. The comparison between the sets of grids for the top-down approach is provided in Table 3a. (The data for one of the experts using the bottom-up approach is incomplete, so a comparison between the two experts was not made there.) The experts only agreed in about 1/3 - 1/2 of their indicated relationships. Furthermore, the final set of relationships for the top-down approach is very different from the final set for the bottom-up approach. (See Table 3b.) When deciding on the relationships during the top-down development, the experts even decided to combine some of the interpretations used in the bottom-up approach, feeling there was too little difference in meaning between them to be significant, and they also dismissed several interpretations during certain time phases (and *tight management* during all time phases) because they felt that the meaning of those interpretations could not be captured by the available metrics in that particular time period. We believe that the differences between the two approaches are mainly due to two facts: (1) the experts were seeing the data from a very different point of view, and (2) the metrics are not ideal, and the experts were not completely certain of all of the relationships that they stated so they changed their opinions over time. However, there were certain relationships which proved more consistent than others. For example, the two experts had strong agreement over the relationships involving programmer hours per line of code, software changes per line of code, and computer time per computer run. These metrics seem to be better understood than the others probably because they are often used for evaluation and comparisons in this field. They also had fairly good agreement with the interpretations of *error prone code*, *lots of reused code*, and *loose management*. The top-down and bottom-up expert

systems had good agreement over programmer hours per line of code and software changes per line of code and over the interpretations of *error prone code* and *good solid code*.

The results of evaluating the four expert systems are displayed in Tables 4.1 and 4.2. The entries in the agreement column are the number of interpretations which were indicated by both the expert system and the information in the database. The first number depicts those interpretations which were explicit in the database, while the second number represents those which were marginally indicated. The entries in the disagreement column are those interpretations indicated by the database, but not listed by the expert system. Again, the first number represents those which are certain and the second number is those which are marginal. Finally, the column labeled "Extra" specifies the number of extra interpretations listed by the expert system. This number is not that meaningful in determining the performance of the rule-based systems at this time because, as discussed previously, the rule-based systems were built to provide as complete a list of interpretations as possible. The manager would then have to decide which interpretations are meaningful and disregard the others. However, in general, it is better to have as few extra interpretations as possible. It should be noted that the total number of interpretations varies from table to table. This is because certain metrics were not available for some projects in some of the time phases. It would be unfair to say the expert systems did not detect certain interpretations if they were not given the manifestations necessary to do so, so these interpretations were not included in the results of the evaluation for those particular cases.

Table 3a - Comparison of Responses Provided by Experts in Each of the Five Time Phases for the Top-Down Expert Systems

Time Phase	Number of Relationships Indicated by Experts		
	Expert 1	Expert 2	Intersection
Early Coding	66	60	23
Middle Coding	78	65	28
Late Coding	81	68	38
Systems Test	79	48	30
Acceptance Test	68	42	23

Table 3b - Comparison of Final Bottom-Up and Final Top-Down Expert Systems

Time Phase	Number of Relationships Used in Each Approach		
	Bottom-Up	Top-Down	Intersection
Early Coding	61	35	15
Middle Coding	65	43	19
Late Coding	63	50	23
Systems Test	65	40	17
Acceptance Test	62	37	17

Table 4.1a - Agreement between Expert System and Information in Database  
Bottom-Up Systems, Early Coding Phase

Project	Rule-Based Deduction System			Frame-Based Abduction System		
	Agreement	Disagreement	Extra	Agreement	Disagreement	Extra
1	1 - 0	0 - 0	2	1 - 0	0 - 0	2
2	0 - 3	3 - 1	7	0 - 0	3 - 4	1
3	0 - 1	0 - 0	9	0 - 1	0 - 0	5
4	0 - 0	3 - 1	0	0 - 0	3 - 1	0
5	2 - 0	2 - 0	5	1 - 0	3 - 0	3
6	0 - 1	1 - 2	3	0 - 1	1 - 2	3
7	1 - 0	3 - 2	1	1 - 0	3 - 2	1
8	0 - 0	3 - 1	0	0 - 0	3 - 1	0
9	1 - 3	2 - 0	8	0 - 0	3 - 3	1
10	5 - 2	2 - 0	4	1 - 0	6 - 2	0
Total	10 - 10	19 - 7	39	4 - 2	25 - 15	16
Percent Agreement	Explicit	34% (10/29)		14% (4/29)		
	Marginal	59% (10/17)		12% (2/17)		
	Combined	43% (20/46)		13% (6/46)		

Table 4.1b - Agreement between Expert System and Information in Database  
Bottom-Up Systems, Middle Coding Phase

Project	Rule-Based Deduction System			Frame-Based Abduction System		
	Agreement	Disagreement	Extra	Agreement	Disagreement	Extra
1	2 - 1	0 - 0	7	0 - 0	2 - 1	1
2	0 - 0	3 - 4	0	0 - 0	3 - 4	0
3	0 - 1	0 - 1	11	0 - 0	0 - 2	1
4	1 - 1	2 - 0	1	1 - 1	2 - 0	1
5	0 - 0	5 - 0	0	0 - 0	5 - 0	0
6	1 - 3	0 - 0	7	0 - 0	1 - 3	1
7	2 - 0	2 - 2	4	0 - 0	4 - 2	2
8	4 - 1	1 - 2	2	1 - 0	4 - 3	0
9	1 - 1	2 - 2	2	1 - 1	2 - 2	2
10	4 - 1	3 - 1	5	1 - 0	6 - 2	3
Total	15 - 9	18 - 12	39	4 - 2	29 - 19	11
Percent Agreement	Explicit	45% (15/33)		12% (4/33)		
	Marginal	43% (9/21)		10% (2/21)		
	Combined	44% (24/54)		11% (6/54)		

Table 4.1c - Agreement between Expert System and Information in Database  
Bottom-Up Systems, Late Coding Phase

Project	Rule-Based Deduction System			Frame-Based Abduction System		
	Agreement	Disagreement	Extra	Agreement	Disagreement	Extra
1	2 - 0	0 - 1	6	0 - 0	2 - 1	2
2	1 - 3	2 - 1	8	1 - 1	2 - 3	4
3	0 - 0	0 - 2	0	0 - 0	0 - 2	0
4	1 - 0	2 - 1	3	1 - 0	2 - 1	0
5	0 - 0	6 - 0	5	0 - 0	6 - 0	5
6	1 - 2	0 - 1	5	0 - 1	1 - 2	1
7	0 - 1	4 - 1	0	0 - 1	4 - 1	0
8	3 - 0	2 - 3	3	1 - 0	4 - 3	2
9	1 - 0	2 - 3	2	1 - 0	2 - 3	1
10	0 - 1	7 - 1	3	0 - 1	7 - 1	3
Total	9 - 7	25 - 14	35	4 - 4	30 - 17	18
Percent Agreement	Explicit	26% (9/34)		12% (4/34)		
	Marginal	33% (7/21)		19% (4/21)		
	Combined	29% (16/55)		15% (8/55)		

Table 4.1d - Agreement between Expert System and Information in Database  
Bottom-Up Systems, Systems Test Phase

Project	Rule-Based Deduction System			Frame-Based Abduction System		
	Agreement	Disagreement	Extra	Agreement	Disagreement	Extra
1	1 - 0	1 - 1	8	1 - 0	1 - 1	6
2	0 - 3	3 - 1	7	0 - 0	3 - 4	1
3	0 - 0	0 - 1	0	0 - 0	0 - 1	0
4	1 - 1	2 - 0	4	1 - 0	2 - 1	0
5	1 - 0	5 - 0	0	1 - 0	5 - 0	0
6	1 - 2	0 - 1	5	0 - 0	1 - 3	1
7	1 - 1	3 - 1	3	1 - 1	3 - 1	3
8	3 - 0	2 - 3	1	3 - 0	2 - 3	1
9	1 - 1	2 - 2	6	1 - 1	2 - 2	6
10	0 - 1	7 - 1	3	0 - 1	7 - 1	3
Total	9 - 9	25 - 11	37	8 - 3	26 - 17	21
Percent Agreement	Explicit	26% (9/34)		24% (8/34)		
	Marginal	45% (9/20)		15% (3/20)		
	Combined	33% (18/54)		20% (11/54)		

Table 4.1e - Agreement between Expert System and Information in Database  
Bottom-Up Systems, Acceptance Test Phase

Project	Rule-Based Deduction System			Frame-Based Abduction System		
	Agreement	Disagreement	Extra	Agreement	Disagreement	Extra
1	1 - 0	3 - 3	2	1 - 0	3 - 3	2
2	0 - 3	3 - 1	9	0 - 2	3 - 2	3
3	0 - 0	0 - 2	4	0 - 0	0 - 2	4
4	0 - 0	3 - 1	0	0 - 0	3 - 1	0
5	3 - 0	3 - 0	5	3 - 0	3 - 0	5
6	1 - 1	0 - 2	2	1 - 1	0 - 2	2
7	3 - 2	1 - 0	3	1 - 0	3 - 2	1
8	1 - 0	4 - 3	1	1 - 0	4 - 3	1
9	1 - 0	2 - 3	2	1 - 0	2 - 3	1
10	0 - 1	7 - 1	3	0 - 1	7 - 1	3
Total	10 - 7	26 - 16	31	8 - 4	28 - 19	22
Percent Agreement	Explicit	28% (10/36)		22% (8/36)		
	Marginal	30% (7/23)		17% (4/23)		
	Combined	29% (17/59)		20% (12/59)		

Table 4.2a - Agreement between Expert System and Information in Database  
Top-Down Systems, Early Coding Phase

Project	Rule-Based Deduction System			Frame-Based Abduction System		
	Agreement	Disagreement	Extra	Agreement	Disagreement	Extra
1	0 - 0	1 - 2	2	0 - 0	1 - 2	2
2	0 - 1	2 - 1	5	0 - 1	2 - 1	2
3	0 - 0	0 - 1	5	0 - 0	0 - 1	1
4	0 - 0	4 - 0	0	0 - 0	4 - 0	0
5	1 - 0	1 - 0	3	0 - 0	2 - 0	2
6	0 - 1	1 - 1	2	0 - 1	1 - 1	2
7	0 - 1	2 - 0	2	0 - 1	2 - 0	2
8	0 - 0	1 - 2	0	0 - 0	1 - 2	0
9	0 - 2	1 - 1	6	0 - 0	1 - 3	2
10	3 - 0	2 - 0	4	2 - 0	3 - 0	1
Total	4 - 5	15 - 8	29	2 - 3	17 - 10	14
Percent Agreement	Explicit	21% (4/19)		11% (2/19)		
	Marginal	38% (5/13)		23% (3/13)		
	Combined	28% (9/32)		16% (5/32)		



Table 4.2b - Agreement between Expert System and Information in Database  
Top-Down Systems, Middle Coding Phase

Project	Rule-Based Deduction System			Frame-Based Abduction System		
	Agreement	Disagreement	Extra	Agreement	Disagreement	Extra
1	1 - 1	1 - 1	6	0 - 0	2 - 2	4
2	0 - 0	2 - 3	0	0 - 0	2 - 3	0
3	0 - 1	0 - 1	7	0 - 1	0 - 1	3
4	0 - 0	2 - 0	1	0 - 0	2 - 0	1
5	0 - 0	4 - 0	0	0 - 0	4 - 0	0
6	1 - 2	0 - 1	6	0 - 0	1 - 3	1
7	0 - 0	2 - 2	4	0 - 0	2 - 2	2
8	2 - 1	2 - 2	0	2 - 0	2 - 3	0
9	0 - 1	2 - 2	3	0 - 1	2 - 2	3
10	5 - 0	1 - 2	3	3 - 0	3 - 2	2
Total	9 - 6	16 - 14	30	5 - 2	20 - 18	16
Percent Agreement	Explicit	36% (9/25)		20% (5/25)		
	Marginal	30% (6/20)		10% (2/20)		
	Combined	33% (15/45)		16% (7/45)		

Table 4.2c - Agreement between Expert System and Information in Database  
Top-Down Systems, Late Coding Phase

Project	Rule-Based Deduction System			Frame-Based Abduction System		
	Agreement	Disagreement	Extra	Agreement	Disagreement	Extra
1	1 - 1	2 - 1	6	0 - 0	3 - 2	4
2	1 - 3	1 - 1	8	0 - 2	2 - 2	1
3	0 - 0	0 - 2	0	0 - 0	0 - 2	0
4	1 - 0	2 - 1	1	1 - 0	2 - 1	1
5	0 - 0	5 - 0	3	0 - 0	5 - 0	3
6	1 - 2	0 - 1	5	0 - 1	1 - 2	2
7	0 - 0	4 - 2	2	0 - 0	4 - 2	2
8	2 - 1	3 - 2	1	1 - 0	4 - 3	0
9	0 - 1	1 - 1	3	0 - 0	1 - 2	1
10	1 - 0	6 - 2	3	0 - 0	7 - 2	1
Total	7 - 8	24 - 13	32	2 - 3	29 - 18	15
Percent Agreement	Explicit	23% (7/31)		6% (2/31)		
	Marginal	38% (8/21)		14% (3/21)		
	Combined	29% (15/52)		10% (5/52)		

Table 4.2d - Agreement between Expert System and Information in Database  
Top-Down Systems, Systems Test Phase

Project	Rule-Based Deduction System			Frame-Based Abduction System		
	Agreement	Disagreement	Extra	Agreement	Disagreement	Extra
1	1 - 1	2 - 1	5	1 - 0	2 - 2	4
2	1 - 3	1 - 0	8	0 - 0	2 - 3	2
3	0 - 0	0 - 0	0	0 - 0	0 - 0	0
4	1 - 0	2 - 1	3	0 - 0	3 - 1	1
5	0 - 0	5 - 0	3	0 - 0	5 - 0	3
6	0 - 1	1 - 1	5	0 - 0	1 - 2	3
7	0 - 1	4 - 1	4	0 - 0	4 - 2	3
8	1 - 0	4 - 3	0	1 - 0	4 - 3	0
9	0 - 1	2 - 1	5	0 - 1	2 - 1	3
10	1 - 0	5 - 2	3	0 - 0	6 - 2	2
Total	5 - 7	26 - 10	36	2 - 1	29 - 16	21
Percent Agreement	Explicit	16% (5/31)		6% (2/31)		
	Marginal	41% (7/17)		6% (1/17)		
	Combined	25% (12/48)		6% (3/48)		

Table 4.2e - Agreement between Expert System and Information in Database  
Top-Down Systems, Acceptance Test Phase

Project	Rule-Based Deduction System			Frame-Based Abduction System		
	Agreement	Disagreement	Extra	Agreement	Disagreement	Extra
1	0 - 0	3 - 2	5	0 - 0	3 - 2	5
2	1 - 2	1 - 0	8	0 - 1	2 - 1	0
3	0 - 0	0 - 0	1	0 - 0	0 - 0	1
4	0 - 0	2 - 1	0	0 - 0	2 - 1	0
5	1 - 0	4 - 0	6	1 - 0	4 - 0	3
6	0 - 0	1 - 1	1	0 - 0	1 - 1	1
7	1 - 1	3 - 1	4	1 - 1	3 - 1	4
8	1 - 1	2 - 2	0	1 - 1	2 - 2	0
9	0 - 0	2 - 1	2	0 - 0	2 - 1	1
10	0 - 0	5 - 2	3	0 - 0	5 - 2	1
Total	4 - 4	23 - 10	30	3 - 3	24 - 11	16
Percent Agreement	Explicit	15% (4/27)		11% (3/27)		
	Marginal	29% (4/14)		21% (3/14)		
	Combined	20% (8/41)		15% (6/41)		

The expert systems performed moderately well given the following limitations: (1) so much of the knowledge and relationships are unclear in this field, (2) the experts themselves do not agree on much of the knowledge, (3) the expert systems used only five

variables and only nine metrics derived from these variables to achieve the list of interpretations, (4) the metrics used are not ideal, (5) many of the interpretations in the database are subjective in nature and therefore may not always be correct, and (6) there may be discrepancies between the interpretations of the particular phase and the overall interpretations for the project.

The systems which were developed with the bottom-up approach performed better than those developed with the top-down approach, and the rule-based deduction systems performed better than the frame-based abduction systems. Both the bottom-up and top-down rule-based systems performed better than either of the frame-based systems. The bottom-up rule-based system performed best, agreeing with an average of 36% (ranging from 29% to 44% depending on time phase) of the combined (both explicit and marginal) interpretations from the database, and the top-down rule-based system agreed with an average of 27% (ranging from 20% to 33%) of the database conclusions. The bottom-up frame-based system agreed with an average of 16% (ranging from 11% to 20%) of the database interpretations, and the top-down frame-based system agreed with an average of 13% (ranging from 6% to 16%) of the database conclusions. Interestingly, both top-down systems and the bottom-up rule-based system agreed with more of the marginal database interpretations than the certain database interpretations. The opposite was true for the bottom-up frame-based system which agreed with an average of 17% of the more certain database conclusions and an average of 15% of the less certain conclusions. It should be pointed out that each expert system produced relatively consistent results throughout its five time phases.

The bottom-up systems contained more relationships between manifestations and interpretations than did the top-down systems. One might assume that the only reason the bottom-up systems agreed with a higher percentage of the database conclusions was that the bottom-up systems would list more interpretations for the same input manifestations (test case). If it listed more interpretations, it would get more right by chance. However, there was not that big a difference between the number of manifestations per interpretation for the bottom-up systems which was 3.16 and the number for the top-down systems which was 2.77. As mentioned before, during the top-down development, the experts combined certain interpretations and dismissed others altogether during certain time phases so there were fewer interpretations for each phase. Although the intent was to throw out inappropriate interpretations and make the top-down systems that much better, the bottom-up systems still captured a higher percentage of correct relationships than did the top-down systems. The total number of interpretations listed by the bottom-up rule-based system was 276 in the 50 test cases. Of these, 95 were in agreement with the database conclusions. The total number of interpretations listed by the top-down rule-based system was 216, and of these, 59 agreed with the database conclusions. Therefore, the bottom-up rule-based system had an average of 34% (95/276) correct interpretations out of all those listed, while the top-down rule-based system averaged only 27% (59/216) correct interpretations.

It is interesting to observe that within both the bottom-up and top-down sets of systems the frame-based system always provided a subset of the interpretations listed by the rule-based system (although in 48% of the combined bottom-up and top-down cases, the rule-based and frame-based systems listed the exact same interpretations). As stated previously, the relationships between the manifestations and interpretations were identical in the frame-based and rule-based systems within each knowledge acquisition approach used. Then, by the nature of the expert system methodologies, the rule-based

system always listed every interpretation associated with every input manifestation, while the frame-based system only provided answers of minimum cardinality which explained all of the manifestations. Since the relationships in the two systems were identical, the frame-based systems could only list the exact same interpretations or a proper subset of those listed by the rule-based systems. As a result, the frame-based systems could not perform better than the rule-based systems with respect to agreement with the database conclusions. The frame-based systems listed an average of 50% fewer extra interpretations (ranging from 29% to 72% depending on time phase) for the bottom-up approach and an average of 48% fewer extra interpretations (ranging from 42% to 53%) for the top-down approach. However, it is better to have extra interpretations than to miss correct interpretations.

The results of using the Kappa statistic to evaluate the expert systems is shown in Table 5. According to these results, the bottom-up rule-based system performed best again, indicating better than chance agreement for more of the interpretations than the other systems did. A few of the interpretations performed relatively well in all or most of the expert systems. These were *low productivity*, *loose management*, *error prone code*, and *computer problems*. The experts had fairly good agreement with each other and also over time (between the bottom-up and the top-down approaches) on the manifestations for *loose management* and *error prone code*. They agreed less on *low productivity* and mostly disagreed on *computer problems*. The interpretations of *low complexity*, *simple system*, and *changes hard to make* also did a little better than chance agreement. The experts had fair agreement with each other and over time concerning *changes hard to make*, but mostly disagreed over *low complexity* and *simple system*. It is interesting to note that the interpretations involving testing performed better in both bottom-up systems than in the top-down systems in general. Perhaps testing is better understood using a very environment-specific approach. Several of the interpretations did not perform well in any of the expert systems, doing worse than chance agreement in all or most cases. These were *high complexity (tough problem)*, *compute bound algorithm*, *good solid code*, *lots of reused code*, *lots of testing*, *little testing*, *lack of thorough testing*, and *tight management*.

## 6. DISCUSSION

The goal of this study was to determine whether it is possible to build useful expert systems for software engineering. Some of the questions which we tried to resolve involved determining how to do the knowledge acquisition and what type of expert system methodology might be best suited for this field.

It was found that the major limitation to developing expert systems for software engineering is the fact that much of the knowledge in this field is not well understood yet. Knowledge was gathered from two experts who have had a great deal of experience in this field, and it was found that they did not agree with each other about many of the relationships we were trying to determine. Furthermore, they did not always agree with themselves when looking at the data from a different point of view at a later date.

Two approaches to knowledge acquisition were used and compared. The bottom-up approach produced better results than the top-down approach. This may well be because the bottom-up approach is more environment-specific. Since the field of

Table 5 - Kappa Statistic Values of Each Interpretation  
in Each of the Four Expert Systems

Interpretation	Bottom-Up Systems		Top-Down Systems	
	RBD	FBA	RBD	FBA
Unstable Specifications	0.120	0.000	-0.065	-0.158
Low Productivity	0.270	-0.065	0.369	0.023
High Productivity	0.000	0.000	0.000	0.000
High Complexity (Tough Problem)	-0.261	-0.236	-0.346	-0.160
Compute Bound Algorithm	-0.139	-0.154	-0.253	-0.168
Low Complexity	0.122	-0.066	0.016	0.155
Simple System	0.121	0.124	***	***
Error Prone Code	0.178	0.118	0.046	0.130
Good Solid Code	-0.134	-0.174	-0.372	-0.082
Lots of Reused Code	-0.121	-0.109	-0.075	-0.163
Lots of Testing	-0.040	0.000	-0.273	-0.205
Little Testing	0.051	-0.144	-0.308	-0.238
Good Testing	0.231	0.296	-0.326	-0.198
Poor Testing	0.186	0.188	-0.241	-0.267
Lack of Thorough Testing	-0.190	-0.061	***	***
Changes Hard to Make	0.000	-0.092	0.211	0.149
Loose Management	0.124	0.123	0.427	0.194
Tight Management	-0.062	-0.114	***	***
Computer Problems	0.235	0.091	0.104	-0.092
Lots of Terminal Jockeys	0.049	-0.087	0.052	0.107

Note -  $K > 0$  indicates better than chance agreement;  $K = 0$  indicates chance agreement;  $K < 0$  indicates worse than chance agreement.

RBD - Rule-Based Deduction; FBA - Frame-Based Abduction

\*\*\* - these interpretations were not used in the top-down systems

software engineering is still new, it is probably better to develop expert systems for one homogeneous environment rather than trying to determine general truths across environments.

The two expert system methodologies, rule-based deduction and frame-based abduction, were also compared with respect to ease of implementation and accuracy of results. The initial knowledge was derived from empirical software engineering research and organized in a table format, so the very first sets of simple rules and frames which were not time-line specific were straightforward to develop. The situation became more complex when the interpretations were made time-line specific. Each frame-based system was divided into five systems based on time period because the second dimension of time could not be incorporated into the frames in a reasonable manner. Furthermore, an attempt was made to rewrite the rules to contain more meaningful and complex relationships among the manifestations in the antecedents. However, it was decided to retain the format of simple rules in order to be as complete as possible. It should be noted that for this type of diagnostic problem in a well-defined domain, it is generally much easier and more natural to write frames than to encode the same information in complex rules

[Ramsey86].

In 48% of the cases, the rule-based and frame-based systems provided the same interpretations. However, when analyzing the results from all projects, the rule-based systems provided more interpretations and exhibited a higher rate of agreement with the database than did the frame-based systems. This is directly attributable to two facts: (1) simple rules were used in the rule-based systems, allowing completeness of the list of interpretations, and (2) the frame-based systems only provide those explanations of minimum cardinality. Therefore, we conclude that the rule-based system with simple rules is probably more applicable to the field of software engineering at this point in time. However, this may very well not be true in the future, as more is learned in this field.

The expert systems performed moderately well, especially when one considers that many of the relationships between the metrics and the interpretations are unclear. The experts did not agree on many of the relationships, and the expert systems cannot perform better than the information included in them. Indeed, the bottom-up rule-based system performed about as well as the experts agreed with each other. In addition, a relatively small number of metrics were used to suggest many interpretations, and the metrics used were not ideal. The experts felt that some of the interpretations could not be adequately described in terms of the available metrics. For example, it was felt that the complexity interpretations could not be adequately captured without error metric data. The experts even threw out one of the interpretations altogether when they were determining relationships using the top-down approach. However, the five variables used in the metrics were easily obtainable, and this is also an important factor when creating expert systems.

Another fact we would like to stress is that the expert systems for the earlier time phases also performed well. This is especially important because a manager should learn of potential problems as early in the development process as possible. Expert systems can be very helpful because they may detect problems which a manager may not recognize early on.

This study has provided many additional new insights into the development of expert systems for software engineering. It is feasible to develop prototype expert systems at this point in time, but we need to define better the relationships that exist between the components. In particular one must define what development characteristics would result in what types of abnormal measures, how this changes through various project development phases, and how certain one is that an abnormal measure results from a certain characteristic. As more is learned in this area, more can be incorporated into useful expert systems.

## 7. FUTURE RESEARCH DIRECTIONS

The development of ARROWSMITH-P was a preliminary attempt at constructing expert systems for software engineering management. The information contained in the knowledge base can be refined, and new knowledge, such as information about error metrics [Weiss85], [Basili84a], can be incorporated into these systems as more is learned. As incorrect relationships are brought to the surface, the systems can be changed to incorporate the knowledge gained from testing. Eventually, the rules should become

more complex as relationships between manifestations and causes become better defined. In addition, the testing of current, ongoing projects can be performed on the expert systems. The data from the new projects can then be incorporated into the environment-specific baselines of metrics so the systems continue to be updated as the environment changes.

In a more general sense, a theoretical framework for developing expert systems for software engineering is needed. For example, a categorization scheme, which would address such issues as when a top-down system is better than a bottom-up system and vice versa, should be built. Also, perhaps a new and different type of inference mechanism or method for building expert systems would better suit the needs of some aspects in this field. All of these issues require a great deal of further research and analysis.

## **8. ACKNOWLEDGEMENT**

The authors are grateful to Frank McGarry, Dr. Jerry Page, Dr. James Reggia, James Ramsey, Bill Decker, and Dave Card for their invaluable assistance in this project. The authors would also like to thank the members of their research group for enlightening comments and ideas.

## References

- [Basili77] Basili, V. R., M. V. Zelkowitz, F. E. McGarry, R. W. Reiter, Jr., W. F. Truskowski, and D. M. Weiss, The Software Engineering Laboratory, SEL-77-001, Software Engineering Laboratory, NASA/Goddard Space Flight Center, Greenbelt, Maryland, May 1977.
- [Basili78] Basili, V. R. and M. V. Zelkowitz, Analyzing Medium-Scale Software Developments, pp. 116-123 in *Proceedings of the Third International Conference on Software Engineering*, Atlanta, Georgia, May 1978.
- [Basili84a] Basili, V. R. and B. T. Perricone, Software Errors and Complexity: An Empirical Investigation, *Communications of the ACM* **27**, 1, pp. 42-52, Jan. 1984.
- [Basili84b] Basili, V. R. and D. M. Weiss, A Methodology for Collecting Valid Software Engineering Data, *IEEE Transactions on Software Engineering* **SE-10**, 6, pp. 728-738, Nov. 1984.
- [Basili85] Basili, V. R. and C. L. Ramsey, ARROWSMITH-P - A Prototype Expert System for Software Engineering Management, pp. 252-264 in *Proceedings of the Expert Systems in Government Symposium*, IEEE, McLean, Virginia, October 1985.
- [Campbell82] Campbell, A. N., V. F. Hollister, R. O. Duda, and P. E. Hart, Recognition of a Hidden Mineral Deposit by an Artificial Intelligence Program, *Science* **217**, pp. 927-928, 3 September 1982.
- [Card82] Card, D. N., F. E. McGarry, J. Page, S. Eslinger, and V. R. Basili, The Software Engineering Laboratory, SEL-81-104, Software Engineering Laboratory, NASA/Goddard Space Flight Center, Greenbelt, Maryland, Feb. 1982.
- [Cohen68] Cohen, J., Weighted Kappa: Nominal Scale Agreement with Provision for Scaled Disagreement or Partial Credit, *Psychological Bulletin* **70**, pp. 213-220, 1968.



[Doerflinger83]

Doerflinger, C. and V. Basili, Monitoring Software Development Through Dynamic Variables, pp. 434-445 in *Proceedings of the IEEE Computer Society's International Computer Software and Applications Conference*, Nov. 1983. (also to appear in IEEE Transactions on Software Engineering).

[Hayes-Roth78]

Hayes-Roth, F., D. Waterman, and D. Lenat, Principles of Pattern-Directed Inference Systems, pp. 577-601 in *Pattern-Directed Inference Systems*, ed. Waterman and Hayes-Roth, Academic Press, 1978.

[Miller82]

Miller, R., H. Pople, and J. Myers, Internist-1: An Experimental Computer-Based Diagnostic Consultant for General Internal Medicine, *New England Journal of Medicine* **307**, pp. 468-476, 1982.

[Minsky75]

Minsky, M., A Framework for Representing Knowledge, pp. 211-277 in *The Psychology of Computer Vision*, ed. P. Winston, McGraw-Hill, Inc., 1975.

[Nau84]

Nau, D. S. and J. A. Reggia, Relationships Between Deductive and Abductive Inference in Knowledge-Based Diagnostic Expert Systems, pp. 500-509 in *Proceedings of the First International Workshop on Expert Database Systems*, 1984.

[Pauker76]

Pauker, S. et al, Towards the Simulation of Clinical Cognition, *American Journal of Medicine* **60**, pp. 981-996, 1976.

[Peng84]

Peng, Y., A General Theoretical Model for Abductive Diagnostic Expert Systems, Tech. Report TR-1402, Computer Science Department, University of Maryland, May 1984.

[Ramsey86]

Ramsey, C., J. Reggia, D. Nau, and A. Ferrentino, A Comparative Analysis of Methods for Expert Systems, *International Journal of Man-Machine Studies*, 1986. Accepted for publication.

[Reggia82a]

Reggia, J. and B. Perricone, KMS Reference Manual, Tech. Report TR-1136, Computer Science Department, University of Maryland, 1982.

[Reggia82b]

Reggia, J., Computer-Assisted Medical Decision Making, pp. 198-213 in *Applications of Computers in Medicine*, ed. M. Schwartz, IEEE Press, 1982.

[Reggia83a]

Reggia, J., D. Nau, and P. Wang, Diagnostic Expert Systems Based on a Set Covering Model, *International Journal of Man-Machine Studies*, pp. 437-460, Nov. 1983.

[Reggia83b]

Reggia, J., D. Nau, and P. Wang, A Theory of Abductive Inference in Diagnostic Expert Systems, Tech. Report TR-1338, Computer Sci. Dept., Univ. of Maryland, College Park, MD, December 1983.

[SEL82]

Annotated Bibliography of Software Engineering Laboratory (SEL) Literature, SEL-82-006, Software Engineering Laboratory, NASA/Goddard Space Flight Center, Greenbelt, Maryland, Nov. 1982.

[Shortliffe76]

Shortliffe, E., *Computer-Based Medical Consultations: MYCIN*, Elsevier, 1976.

[Shubin82]

Shubin, H. and J. Ulrich, IDT: An Intelligent Diagnostic Tool, pp. 290-295 in *Proceedings of the National Conference on Artificial Intelligence, AAAI*, 1982.

[Spitzer67]

Spitzer, R., J. Cohen, J. Fleiss, and J. Endicott, Quantification of Agreement in Psychiatric Diagnosis, *Archives of General Psychiatry* 17, pp. 83-87, 1967.

[Weiss85]

Weiss, D. M. and V. R. Basili, Evaluating Software Development by Analysis of Changes: Some Data From the Software Engineering Laboratory, *IEEE Transactions on Software Engineering* SE-11, 2, pp. 157-168, Feb. 1985.

APPENDIX 1a - A sample interactive session with the rule-based deduction expert system.

THIS EXPERT SYSTEM WILL HELP A MANAGER OF A SOFTWARE PROJECT  
DETERMINE IF THE PROJECT IS ON SCHEDULE OR IN TROUBLE.  
PLEASE ANSWER THE FOLLOWING QUESTIONS.

COMPUTER RUNS PER LINE OF SOURCE CODE:

- (1) ABOVE NORMAL
  - (2) NORMAL
  - (3) BELOW NORMAL
- = ?

2.

COMPUTER TIME PER LINE OF SOURCE CODE:

- (1) ABOVE NORMAL
  - (2) NORMAL
  - (3) BELOW NORMAL
- = ?

2.

SOFTWARE CHANGES PER LINE OF SOURCE CODE:

- (1) ABOVE NORMAL
  - (2) NORMAL
  - (3) BELOW NORMAL
- = ?

2.

PROGRAMMER HOURS PER LINE OF SOURCE CODE:

- (1) ABOVE NORMAL
  - (2) NORMAL
  - (3) BELOW NORMAL
- = ?

2.

COMPUTER TIME PER COMPUTER RUN:

- (1) ABOVE NORMAL
  - (2) NORMAL
  - (3) BELOW NORMAL
- = ?

2.

SOFTWARE CHANGES PER COMPUTER RUN:

- (1) ABOVE NORMAL
  - (2) NORMAL
  - (3) BELOW NORMAL
- = ?

2.

PROGRAMMER HOURS PER COMPUTER RUN:

- (1) ABOVE NORMAL
  - (2) NORMAL
  - (3) BELOW NORMAL
- = ?

3.

PROJECT TIME PHASE:

- (1) EARLY CODE PHASE
  - (2) MIDDLE CODE PHASE
  - (3) LATE CODE PHASE
  - (4) SYSTEMS TEST PHASE
  - (5) ACCEPTANCE TEST PHASE
- = ?

2.

COMPUTER TIME PER SOFTWARE CHANGE:

- (1) ABOVE NORMAL

- (2)NORMAL
- (3)BELOW NORMAL
- = ?

2.

PROGRAMMER HOURS PER SOFTWARE CHANGE:

- (1)ABOVE NORMAL
- (2)NORMAL
- (3)BELOW NORMAL
- = ?

3.

POSSIBLE INTERPRETATIONS ARE:

- ERROR PRONE CODE <0.94>
- EASY ERRORS OR CHANGES BEING FOUND OR FIXED <0.81>
- LOTS OF TESTING <0.75>
- LOTS OF TERMINAL JOCKEYS <0.75>
- UNSTABLE SPECIFICATIONS <0.50>
- NEAR BUILD OR MILESTONE DATE <0.50>
- GOOD TESTING OR GOOD TEST PLAN <0.25>
- MODIFICATIONS BEING MADE TO RECENTLY TRANSPORTED CODE <0.25>

Note - User answers are in boldface.

APPENDIX 1b - A sample interactive session with the frame-based abduction expert system.

THIS EXPERT SYSTEM WILL HELP A MANAGER OF A SOFTWARE PROJECT  
DETERMINE IF THE PROJECT IS ON SCHEDULE OR IN TROUBLE.  
THIS PARTICULAR SYSTEM SHOULD BE USED FOR THE MIDDLE CODING PHASE.  
PLEASE ANSWER THE FOLLOWING QUESTIONS.

FOCUS OF SUBPROBLEM:

\*\*\*THIS SUBPROBLEM IS CURRENTLY ACTIVE\*\*\*

GENERATOR:

COMPETING POSSIBILITIES:

UNSTABLE SPECIFICATIONS  
LATE DESIGN  
NEW OR LATE DEVELOPMENT  
LOW PRODUCTIVITY  
HIGH PRODUCTIVITY  
HIGH COMPLEXITY OR TOUGH PROBLEM  
HIGH COMP OR COMPUTE BOUND ALGORITHMS RUN OR TESTED  
LOW COMPLEXITY  
SIMPLE SYSTEM  
REMOVAL OF CODE BY TESTING OR TRANSPORTING  
INFLUX OF TRANSPORTED CODE  
LITTLE EXECUTABLE CODE BEING DEVELOPED  
ERROR PRONE CODE  
GOOD SOLID AND RELIABLE CODE  
NEAR BUILD OR MILESTONE DATE  
LARGE PORTION OF REUSED CODE OR EARLY AND LARGER TESTS  
LOTS OF TESTING  
LITTLE OR NOT ENOUGH ONLINE TESTING BEING DONE  
GOOD TESTING OR GOOD TEST PLAN  
UNIT TESTING BEING DONE  
LACK OF THOROUGH TESTING  
POOR TESTING PROGRAM  
SYSTEM AND INTEGRATION TESTING STARTED EARLY  
CHANGE BACKLOG OR HOLDING CHANGES  
CHANGE BACKLOG OR HOLDING CODE  
CHANGES HARD TO ISOLATE  
CHANGES HARD TO MAKE  
EASY ERRORS OR CHANGES BEING FOUND OR FIXED  
MODIFICATIONS BEING MADE TO RECENTLY TRANSPORTED CODE  
LOOSE CONFIGURATION MANAGEMENT OR UNSTRUCTURED DEV  
TIGHT MANAGEMENT PLAN OR GOOD CONFIGURATION CONTROL  
COMPUTER PROBLEMS OR INACCESSIBILITY OR ENV CONSTRAINTS  
LOTS OF TERMINAL JOCKEYS

COMPUTER RUNS PER LINE OF SOURCE CODE:

- (1) ABOVE NORMAL
  - (2) NORMAL
  - (3) BELOW NORMAL
- = ?

2.

COMPUTER TIME PER LINE OF SOURCE CODE:

- (1) ABOVE NORMAL
  - (2) NORMAL
  - (3) BELOW NORMAL
- = ?

2.

SOFTWARE CHANGES PER LINE OF SOURCE CODE:

- (1) ABOVE NORMAL
  - (2) NORMAL
  - (3) BELOW NORMAL
- = ?

2.

PROGRAMMER HOURS PER LINE OF SOURCE CODE:

- (1) ABOVE NORMAL
- (2) NORMAL
- (3) BELOW NORMAL
- = ?

2.

SOFTWARE CHANGES PER COMPUTER RUN:

- (1) ABOVE NORMAL
- (2) NORMAL
- (3) BELOW NORMAL
- = ?

2.

COMPUTER TIME PER COMPUTER RUN:

- (1) ABOVE NORMAL
- (2) NORMAL
- (3) BELOW NORMAL
- = ?

2.

PROGRAMMER HOURS PER COMPUTER RUN:

- (1) ABOVE NORMAL
- (2) NORMAL
- (3) BELOW NORMAL
- = ?

3.

FOCUS OF SUBPROBLEM:

GENERATOR:

COMPETING POSSIBILITIES:

- LOTS OF TERMINAL JOCKEYS
- EASY ERRORS OR CHANGES BEING FOUND OR FIXED
- LOTS OF TESTING
- ERROR PRONE CODE
- UNSTABLE SPECIFICATIONS

PROGRAMMER HOURS PER SOFTWARE CHANGE:

- (1) ABOVE NORMAL
- (2) NORMAL
- (3) BELOW NORMAL
- = ?

3.

FOCUS OF SUBPROBLEM:

GENERATOR:

COMPETING POSSIBILITIES:

- EASY ERRORS OR CHANGES BEING FOUND OR FIXED
- ERROR PRONE CODE

COMPUTER TIME PER SOFTWARE CHANGE:

- (1) ABOVE NORMAL
- (2) NORMAL
- (3) BELOW NORMAL
- = ?

2.

POSSIBLE INTERPRETATIONS ARE:

- EASY ERRORS OR CHANGES BEING FOUND OR FIXED <H>
- ERROR PRONE CODE <L>

Note - User answers are in boldface.

- Both interpretations listed as solutions can explain all of the manifestations, but the first is given a high measure of likelihood (shown by the <H>) of being correct, while Error Prone Code is rated low.

**REPORT DOCUMENTATION PAGE**

1a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED		1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION/AVAILABILITY OF REPORT Distribution unlimited	
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE			
4. PERFORMING ORGANIZATION REPORT NUMBER(S) TR-1708		5. MONITORING ORGANIZATION REPORT NUMBER(S)	
6a. NAME OF PERFORMING ORGANIZATION University of Maryland	6b. OFFICE SYMBOL (if applicable) CMSC	7a. NAME OF MONITORING ORGANIZATION	
6c. ADDRESS (City, State, and ZIP Code) Department of Computer Science College Park, MD 20742		7b. ADDRESS (City, State, and ZIP Code)	
8a. NAME OF FUNDING / SPONSORING ORGANIZATION NASA	8b. OFFICE SYMBOL (if applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER	
8c. ADDRESS (City, State, and ZIP Code) Greenbelt, MD 20771		10. SOURCE OF FUNDING NUMBERS	
		PROGRAM ELEMENT NO.	PROJECT NO.
		TASK NO.	WORK UNIT ACCESSION NO.
11. TITLE (Include Security Classification)   Unclassified AN EVALUATION OF EXPERT SYSTEMS FOR SOFTWARE ENGINEERING MANAGEMENT			
12. PERSONAL AUTHOR(S)   Connie Loggia Ramsey and Victor R. Basili			
13a. TYPE OF REPORT	13b. TIME COVERED FROM _____ TO _____	14. DATE OF REPORT (Year, Month, Day) September 1986	15. PAGE COUNT 28
16. SUPPLEMENTARY NOTATION			
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	
19. ABSTRACT (Continue on reverse if necessary and identify by block number) Although the field of software engineering is relatively new, it can benefit from the use of expert systems. Four prototype expert systems have been developed to aid in software engineering management. Given the values for certain metrics, these systems will provide interpretations which explain any abnormal patterns of these values during the development of a software project. The four expert systems, which solve the same problem, were built using two different approaches to knowledge acquisition, a bottom-up approach and a top-down approach, and two different expert system methods, rule-based deduction and frame-based abduction. A comparison was performed to see which methods better suit the needs of this field. It was found that the bottom-up approach led to better results than did the top-down approach, and the rule-based deduction systems using simple rules provided more complete and correct solutions than did the frame-based abduction systems.			
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS		21. ABSTRACT SECURITY CLASSIFICATION	
22a. NAME OF RESPONSIBLE INDIVIDUAL Connie Loggia Ramsey		22b. TELEPHONE (Include Area Code) 301-454-6154	22c. OFFICE SYMBOL CMSC

STATE OF CALIFORNIA

DEPARTMENT OF REVENUE

1917

IN SENATE, January 15, 1917.

REPORT OF THE COMMISSIONER OF REVENUE

FOR THE YEAR ENDING DECEMBER 31, 1916.

ALBINO W. BROWN, COMMISSIONER.

FRANCIS J. WOOD, CLERK.

FRANCIS J. WOOD, CLERK.

FRANCIS J. WOOD, CLERK.

FRANCIS J. WOOD, CLERK.

FRANCIS J. WOOD, CLERK.

FRANCIS J. WOOD, CLERK.

FRANCIS J. WOOD, CLERK.

FRANCIS J. WOOD, CLERK.

FRANCIS J. WOOD, CLERK.

FRANCIS J. WOOD, CLERK.

FRANCIS J. WOOD, CLERK.

FRANCIS J. WOOD, CLERK.