


GUIDON-WATCH: A Graphic Interface for Viewing a Knowledge-Based System

Mark H. Richer
William J. Clancey
Stanford University



Multiple windows, menus, and a mouse help students browse through a complex medical consulting database and view reasoning processes during a consultation.

An increasing number of Artificial Intelligence (AI) programs are implemented on high-performance workstations with a bit-map display, a mouse, and a keyboard. The programming environment (usually a dialect of Lisp) generally provides support for displaying multiple windows and using menus that can be selected with a mouse. Importantly, a programmer can also specify arbitrary regions of a window (e.g., text items) to be selectable with the mouse. This means that a user can invoke an action by pressing or releasing a mouse button while the mouse cursor is in a selectable region. These features make it

GUIDON-WATCH is a graphic interface that uses multiple windows and a mouse to allow a student to browse through a knowledge base and view reasoning processes during diagnostic problem solving. This article presents methods for providing multiple views of hierarchical structures, overlaying results of a search process on top of static structures to make the strategy visible, and graphically expressing evidence relations between findings and hypotheses. We demonstrate the advantages of stating a diagnostic search procedure in a well-structured, rule-based language, separate from domain knowledge. A number of issues in software design are also considered, including the automatic management of a multiple-window display.

possible to create a user interface that is efficient and easy to use for viewing and browsing a complex system.

The GUIDON project at Stanford University is investigating ways in which knowledge-based systems can provide the basis for teaching programs. NEOMYCIN, a medical consultation system, has been developed for this purpose.¹ This article describes GUIDON-WATCH, a graphic interface to NEOMYCIN that uses multiple windows and the mouse to allow a user to browse through the NEOMYCIN knowledge base and view reasoning processes during a consultation. The results reported include methods for providing multiple views of a database, techniques for illustrating dynamic processes including a search strategy, and some conclusions regarding automatic management of a multiple window display.

Project goals

The capability of displaying and selecting information in several windows allows people to control and observe the behavior of an application program in an easy fashion. A graphic interface to a knowledge-based system can serve different kinds of users, including system designers, implementers, domain experts, students, and other end users.

In the GUIDON project, the end users will be medical students. We are currently collaborating with physicians

and medical students to adapt NEOMYCIN, GUIDON-WATCH, and other programs for medical instruction. However, when this work began better tools were also needed to maintain the NEOMYCIN knowledge base and to debug program behavior. As a result, GUIDON-WATCH evolved into a tool for both programmers and medical students to use. We are just starting to make a clean separation between the functionality that is useful for students and that for programmers. We plan to develop user profiles that determine the interface behavior in a given situation. The current prototype can be customized only by making changes at the programming level.

GUIDON-WATCH is based on established principles for designing user interfaces on graphic workstations.²⁻⁵ The design criteria for GUIDON-WATCH emerged from the conventional wisdom on the subject. The user interface is viewed as a conversation consisting of two languages⁵: (1) the language in which the user retrieves or requests information (with the mouse), and (2) the program's display and its interpretation. We aim to maximize expressiveness, understandability, and efficiency for both languages. The user should be able to retrieve all information through one interface that is easy to understand and efficient to use. The display should include all relevant information, be easy to interpret, and quick to update when a user makes a request.

Several GUIDON-WATCH users have found the interface simple, consistent, and easy to use. However, those unfamiliar with NEOMYCIN have difficulty realizing exactly how and when the display can be useful. We have found that the display is the best means we have for explaining NEOMYCIN. Therefore, an on-line introduction to GUIDON-WATCH and NEOMYCIN is planned.

Informal evaluation with Stanford University medical students is scheduled for the fall of 1985. Students will




Figure 1: The evolution of a knowledge-based system. MYCIN evolved into EMYCIN, a domain-independent shell for building knowledge-based systems. The GUIDON tutoring system is a separate module that could be used with any EMYCIN system. EMYCIN was not found to be an adequate foundation for an instructional program. Therefore, EMYCIN and the MYCIN knowledge base were reconfigured into NEOMYCIN, a medical consultation system designed for enhanced explanations and tutoring capabilities. The domain-independent shell that NEOMYCIN is built with is called HERACLES. NEOMYCIN is the basis for GUIDON-2, a tutoring system now in development.

watch NEOMYCIN diagnose one or more patients. Data records of actual patients will be stored in files that can be accessed by NEOMYCIN during its questioning phase. A student will use GUIDON-WATCH to observe NEOMYCIN's reasoning processes during the consultation. NEOMYCIN will be able to explain in English why it asked a question.⁶ Eventually, students will assist NEOMYCIN during a diagnosis in an apprenticeship setting.

The major results to date are summarized here:

1. Multiple windows can provide several concurrent views of a knowledge-based system. They help users cope with the complexity of the system by highlighting and summarizing important reasoning events during a problem-solving session.

2. Several methods for highlighting facts and events were found effective. These include using different font styles, reverse video, boxing, flashing, and graying regions. With these techniques, dynamic information associated with a given patient can be overlaid on top of static structures such as a disorder tree or a table of evidence.

3. Early results indicate that both programmers and medical users prefer to have GUIDON-WATCH manage screen space automatically. This includes the sizing, placing, and closing of windows. It is not trivial to do this with a large number of windows, particularly during development when changes to the system are frequent. A knowledge-based approach to window management is suggested.

The development of NEOMYCIN

The GUIDON project evolved from the MYCIN experiments of the 1970's (Figure 1).^{7,8} MYCIN is a rule-based consultation program that recommends drug therapy for certain infectious diseases (e.g., meningitis). Because much of the functionality (e.g., the inference mechanism) of MYCIN does not depend on medical knowledge, it was possible to develop a domain-independent shell called EMYCIN.⁹ MYCIN now consists of EMYCIN plus the MYCIN medical knowledge base. EMYCIN was used to develop several other knowledge-based systems and is the basis for several commercial products.

In 1979 Clancey completed GUIDON,¹⁰ an intelligent tutoring system that interfaces with EMYCIN. In theory, GUIDON can teach a student the rules in an EMYCIN knowledge base. However, Clancey¹ found that the MYCIN rules were often difficult to understand because they combine a diagnostic procedure with medical facts in an opaque manner. In a MYCIN rule the ordering of conjunct clauses in the premise might *implicitly* contain a strategy. For example, a rule might apply only if the patient is an alcoholic. One MYCIN rule premise begins with, "If the patient is over 18 years of age and an alcoholic." The strategy that is implicitly represented in this rule premise is "Don't ask a patient under 18 years of age if they are alcoholic."

GUIDON demonstrated that satisfying the requirements for expert performance is not necessarily sufficient for the purpose of explanation and tutoring. Therefore, MYCIN was significantly reconfigured into a new program called NEOMYCIN¹ that represents a diagnostic strategy separately from medical facts. For example, the diagnostic strategy used in NEOMYCIN explicitly states that the program should check for conditions that would make a question inappropriate. The knowledge base has also been expanded to include diseases that can be confused with meningitis (this is important for instruction). NEOMYCIN is the foundation for GUIDON-2, a new series of instructional programs. GUIDON-WATCH is the first component of the GUIDON-2 system. Importantly, interactive graphics makes knowledge and reasoning visible only to the extent that the knowledge is represented *explicitly* in a program. The well-structured and explicit design of NEOMYCIN provides many opportunities for exposing the program's reasoning to students and other users.

NEOMYCIN has led to a domain-independent system called HERACLES. HERACLES is to NEOMYCIN as EMYCIN is to MYCIN. In other words, NEOMYCIN consists of HERACLES and a medical knowledge base (Figure 2). HERACLES is a software tool applicable to diagnostic problems in many domains. For example, HERACLES was used to develop a knowledge base for cast-iron fault diagnosis.¹¹ The HERACLES program includes a diagnostic procedure represented in a rule-based declarative language, rule interpreters, a set of domain relations (e.g., causes, subtype, suggests), various software tools for developing knowledge-based systems (many derived from EMYCIN), an explanation facility, and GUIDON-WATCH. To construct a specific consultation program, the system designer adds a knowledge base of facts and rules. All the examples in this article use the NEOMYCIN knowledge base, but GUIDON-WATCH will work with any HERACLES knowledge base.

Description of the GUIDON-WATCH display

This section details the windows and menus used in GUIDON-WATCH. First, the programming environment is briefly described to show the tools we used when we began the project. Then, an overview of the interface is provided. The third section describes the window display facilities in detail.

Programming environment. GUIDON-WATCH is implemented on Xerox 1100 Series workstations running Interlisp-D. The black-and-white display screen is 1024 pixels wide by 808 pixels high, which provides approximately 75 pixels per inch of resolution. Interlisp-D provides a window package that supports multiple overlapping windows, scroll bars, and other window operations.¹² Many graphics primitives are provided for drawing lines and curves, manipulating bit maps, filling and manipulating

regions, checking the state and position of the mouse, and so on. In addition, a menu package, a package to display trees, and several default window functions (e.g., scrolling by repainting) are provided. It required only a page of code to implement a simple pull-down menu package using the window primitives.

Overview of the user interface. Pull-down menus and a Prompt window appear at the top of the GUIDON-WATCH display (Figure 3). The Prompt window is a standard part of the Interlisp-D user interface that is used to print messages. Currently there are three pull-down menus of interest to a medical student: *KB (Knowledge Base) Windows*, *Consult*, and *Help*. The KB Windows pull-down menu displays a list of windows that can be opened for browsing through the knowledge base or viewing a consultation. The Consult menu is used to start and quit a consultation. The Help menu allows the user to obtain information on the contents of a window.

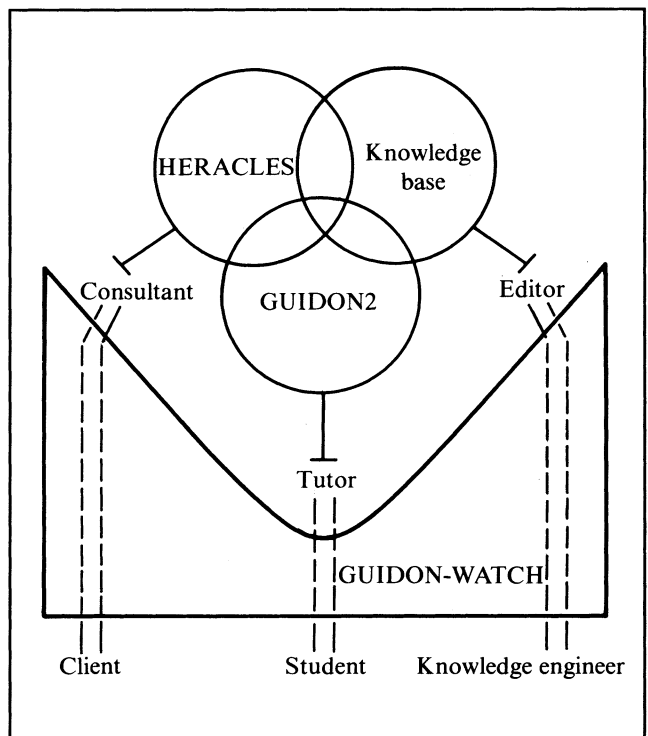


Figure 2: The HERACLES architecture. The relation between GUIDON-WATCH and three primary system modules is illustrated here. The consultation system consists of HERACLES, a knowledge base, and GUIDON-WATCH. For instructional use, the GUIDON-2 module (now in development) can be added. GUIDON-WATCH provides an interface for instructional use, to run consultations and to edit the knowledge base. Although there are differences in the interface for each type of user, in general, the interface is very similar and represents a single program with several modes of behavior. (The graphic editor interface is not described in this paper because it has not been completely integrated with HERACLES.)

Use of the mouse. Xerox 1100 computers can be used with either a two- or three-button mouse (selecting the left and right button at the same time on a two-button mouse is

equivalent to pressing the middle button on the three-button mouse). In GUIDON-WATCH, the mouse is used in a simple and consistent way. The left button is used to

KB Windows | Consult | Help

- Findings
- Hypotheses
- Rules
- Tasks
- Relations
- Causal Association Network
- MetaStrategy
- Taxonomy**
- Differential
- Hypotheses-With-Evidence
- Positive Findings
- TaskStack
- Dynamic Task Tree
- Task History

1) ** Mary 42 YEARS FEMALE

2) Please describe the chief complaints:
 ** HEADACHE
 ** STIFF-NECK-ON-FLEXION
 ** NAUSEA
 **

3) For how long has Mary's headache lasted?
 ** 10 DAYS

4) How severe is Mary's headache (on a scale of 0 to 4 with 0 for very mild and 4 for very severe)?
 ** 3

DIFFERENTIAL:
 (VIRAL-MENINGITIS 200) (CHRONIC-MENINGITIS 200)

5) Does Mary have a fever?
 **

Prompt Window Copyright (C) by Xerox Corporation 19-Mar-85 14:39:37

Continue Till Next? Continue Till a Certain?
 Continue Thru Last? Userexec
 Pageheight 0 Resize
 Quit Explain

Explanation Window

WHYPRUNE
 [i.e. WHY are we asking whether Mary has a fever?]
 [4.0] We are trying to decide whether Mary has infection.
 Whether Mary has a fever is strongly associated with infection.
 ** WHYPRUNE
 [i.e. WHY are we trying to decide whether Mary has infection?]
 [5.0] We are trying to get a general idea of the problem: categorize it into one of several pathogenetic classes or disease locus, or both.

Evidence for INFECTIOUS-PROCESS

FINDING	RULE(S)	MAXCF	MINCF
FEBRILE	RULE423	700	
WBC	RULE350	500	
PMNS	RULE350	500	
BANDS	RULE350	500	

RULE423

If: The patient has a fever
 Then: If you are considering any-disorder, there is suggestive evidence (.7) that the underlying etiology of the patient's illness is infectious-process

Differential

HYPOTHESIS	CF	CUMCF
CHRONIC-MENINGITIS	200	
VIRAL-MENINGITIS	200	

Hypotheses With Evidence

HYPOTHESIS	CF	CUMCF
MENINGITIS	500	500
INFECTIOUS-PROCESS	---	500
CHRONIC-MENINGITIS	200	
VIRAL-MENINGITIS	200	
ACUTE-MENINGITIS	---	200

Figure 3: A GUIDON-WATCH display during consultation. The user is running a Consult, and the system has paused at question 5. The user has opened several windows to get information about hypotheses being considered at this time. The use of pull-down menus is also illustrated. The user has selected the KB windows and moved the mouse over the menu item Taxonomy. If the user releases the mouse button now, the Taxonomy windows will be displayed.

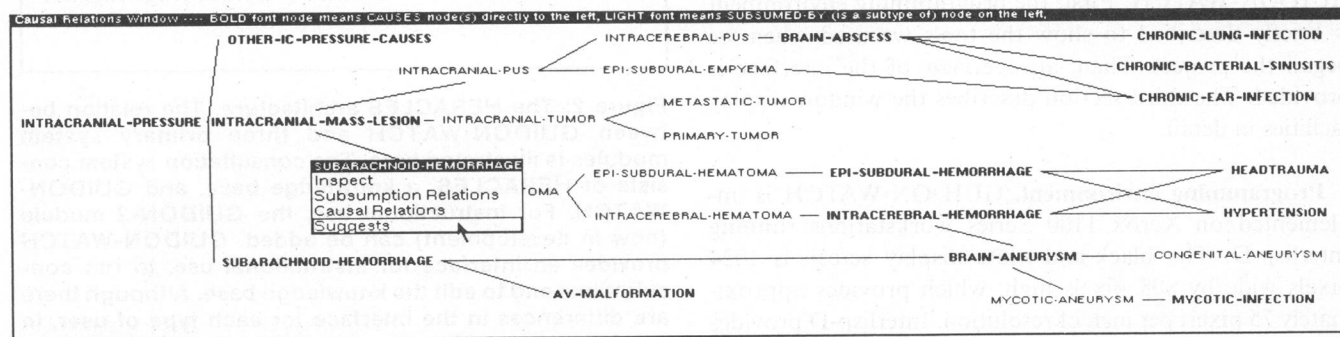


Figure 4: The Causal Relations window. The user has selected the node SUBARACHNOID-HEMORRHAGE with the left button and a pop-up menu displays options for additional information. This graph was automatically generated from the NEOMYCIN knowledge base, edited by hand to fit on the screen, and then stored on a file. If the user wishes to display the graph with a different root node, GUIDON-WATCH dynamically generates the graph at runtime.

select all menu items and text items in a window. For example, in a window that displays a list of diseases, the user can select the name of a disease using the left button. A pop-up menu is displayed that allows a user to get more information in another window (Figure 4). Only those items that are currently relevant appear in the pop-up menu.

It has not been decided whether or not students will be asked to use more than the left button. In our current programming environment, the right button is used in the default manner provided by Interlisp-D, to manipulate windows (e.g., reshaping, closing). The middle button is sometimes used to display a pop-up menu with items that apply to the entire data structure in a window. For example, a user may want to highlight those items in a window that have a certain property. We are considering the use of icons in a window for operations besides selecting menu or text items (e.g., closing a window). Therefore, the student interface may use only one button.

On-line help. If the user selects Help window from the Help pull-down menu, a Help icon attaches to the mouse cursor. The user can get help about a window by moving the Help icon into a window and buttoning the window. A message associated with the selected window is printed in a special help window.

Management of windows. The Interlisp-D graphics package provides functions for prompting users to position a ghost image of a window or to a shape of a window. These prompts can be confusing to novices and distract them from the task at hand. If it is possible to make a good decision regarding the size and position of a window, we can free the user from this chore. In addition, an automatic window-management system can often optimize the use of screen space better than a user can. This is true in GUIDON-WATCH because there are a known set of windows whose contents are constrained to a certain form (e.g., a table).

To manage the window display, we divide the screen conceptually into logical units. The GUIDON-WATCH screen currently consists of top, middle, and bottom sections. The top section contains the pull-down menus and the Prompt window. The bottom and middle sections display knowledge-base structures and have well-defined lower borders. Another logical division of the GUIDON-WATCH display provides vertical boundaries. For example, the width of the screen can be divided into equal or unequal regions. The current prototype uses three regions with two equal and one slightly wider than the other two. Furthermore, you can have a hierarchy of subdivisions (i.e., regions). Each window in GUIDON-WATCH is associated with one or more regions where it can be displayed.

GUIDON-WATCH decides where to place a window based on several considerations: (1) the default region of the window, (2) the other windows that are displayed and

their position, and (3) the set of windows that the user would most likely prefer to remain in view. While the current window management is effective, we would like to extend the flexibility of the interface. This would require a more complex scheme. It might be necessary to consider moving or reshaping windows that are already on the screen. Note that window systems that provide this capability do not consider the *semantics* of the contents of windows. Therefore, algorithms for scaling pictures and changing the font size of text are not adequate when the system must decide where windows should be placed and which windows should be closed or covered.

Although the user relinquishes flexibility and control, the benefits of automatic screen management seem to outweigh potential disadvantages. Automatic window management saves the user time and maximizes the use of screen space. It is possible to allow the user to turn off automatic features, change defaults, or allow the user to use the move and reshape facilities. Furthermore, menus or icons can be used to allow the user to choose from a predefined set of sizes, positions, fonts, and so on, but then the implementation of the automatic window manager becomes increasingly complex. In our current implementation, when a window is displayed, a complex conditional in the window's display function is evaluated. This code is difficult to understand and modify. In addition, the situation is complicated by the need for different user profiles. We are considering an approach where the behavior of the interface is specified separately and declaratively using knowledge-representation formalisms (e.g., rules) and object-oriented programming.

Dynamic updating of the screen display. Displaying dynamically changing information presents problems that are not unique to our application. For example, how often do you update the screen? Do you gray out regions that are out of date or do you update them? Our philosophy is that users should be able to open and close windows at any time and that the display should accurately reflect the current state of the system or gray-out regions that are not continuously updated. Regions that are grayed out can either be updated automatically at specified intervals or manually updated by the user's simply buttoning the window to redisplay itself.

The GUIDON-WATCH windows. Here we describe many of the windows available to the GUIDON-WATCH user and address important issues. What information in a HERACLES knowledge base is most important to display for programmers and for medical students? How can dynamic information be displayed? In the next section we focus on static knowledge structures and the way they are displayed in GUIDON-WATCH. Subsequent sections discuss the display of dynamic consultation knowledge.

What is there to display in a knowledge base? A HERACLES knowledge base (e.g., the NEOMYCIN

medical knowledge base) includes findings, hypotheses, rules, tasks, and relations. Findings are data that can be requested or inferred from rules. Generally findings can be observed or measured. Hypotheses can only be inferred from rules. In NEOMYCIN, hypotheses include diseases and pathophysiological states. Relations refer to predicate calculus relations and in HERACLES include subtype, causes, etc.

Static knowledge includes facts about findings and hypotheses as defined by relations (e.g., meningitis is a subtype of infection, headache is a finding, and so on). It also includes the diagnostic procedure and domain rules (e.g., if the patient has double vision, there is suggestive evidence for intracranial pressure). *Dynamic* knowledge is situation specific and refers to information that becomes known only during a problem-solving session (e.g., Mary's temperature is 102 degrees).

NEOMYCIN uses a diagnostic strategy known as heuristic classification problem solving.¹³ Given an enumerated

set of solutions (e.g., diseases or possible diagnoses), NEOMYCIN heuristically maps a set of findings onto one or more possible solutions. This diagnostic procedure is provided by HERACLES (or Heuristic Classification Shell). It is represented as tasks, which are procedures that are stated in a declarative rule-based language (Figure 5). When a task is invoked, one or more of its metarules are applied (Figures 6 and 7). Metarules in HERACLES are similar to conditionals in a procedure, but they are expressed as abstract rules.

Windows that display static knowledge include the task, metarules, and rule windows in Figures 5, 6, and 7. They also include the Findings, Hypotheses, and Relations windows, which simply display an alphabetical ordering. Other windows display a graph to show the relationships between groups of objects. The Causal Relations window (Figure 4) is a lattice with causal and subtype links between findings and hypotheses; the Diagnostic Strategy window (Figure 8) shows the calling structure of the diagnostic tasks; and the Taxonomy window (Figure 9) represents a subtype hierarchy of disorders. In all of these windows the user may select an item to get more information.

The Taxonomy and Causal Relations windows. An important concept in medical diagnosis is the *differential*, the set of competing hypotheses currently being considered. The etiological taxonomy is a tree of possible diagnoses or solutions in NEOMYCIN. The differential represents a *cut* through this solution space. Boxing the hypotheses in the Taxonomy window that are on the differential is a simple way to make this *cut* visible (Figure 9).

Flashing and boxing nodes in the Taxonomy and Causal Relations windows emphasize the dynamic search strategy. Whenever a hypothesis is added to the differential, its corresponding node label is flashed and then boxed. Whenever the hypothesis is removed from the differential, the box is redrawn with lighter lines, so that the hypotheses that had been considered previously are still highlighted, but the ones currently on the differential are more prominent (Figure 9). A student can observe NEOMYCIN looking *up* the disorder tree to group and compare categories of disorders before looking *down* to refine hypotheses. In essence, we are *reifying* (i.e., making more concrete) the process of problem solving.

Conclusions in a HERACLES consultation are associated with *certainty factors* that represent a degree of belief. They are not probabilities. In HERACLES, each hypothesis has both a certainty factor (CF) and a cumulative certainty factor (CUMCF). The CF represents the combined certainty of all rules that have concluded directly about the hypothesis. The CUMCF represents a combination of the CF of a given hypothesis (which may be zero) with CFs of its descendants in the disorder taxonomy. For example, evidence for meningitis (a positive CF) increases the CUMCF of infectious process because meningitis is a subtype of infectious process. (To be exact, negative CFs of

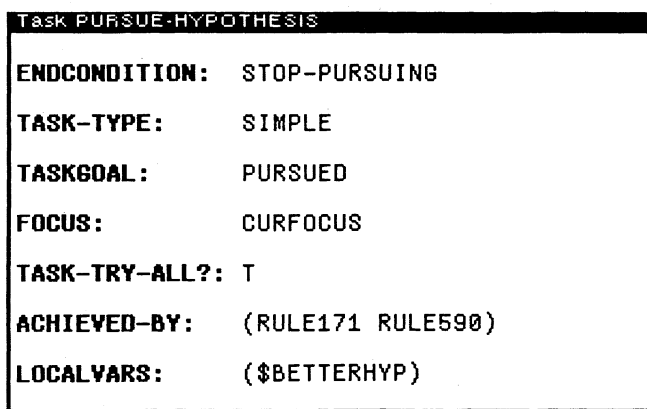


Figure 5: The Task Property window. Here the properties and values of the task Pursue-Hypothesis are displayed.

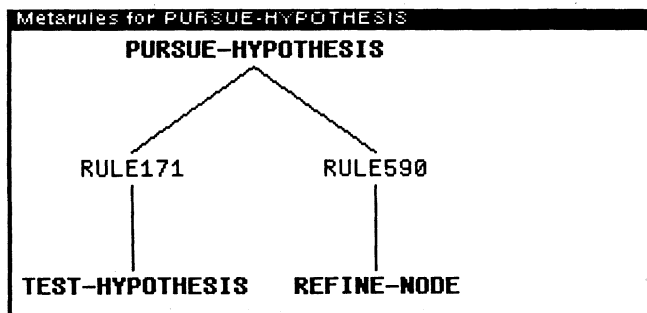


Figure 6: The Metarules window. Here the metarules that the task Pursue-Hypothesis calls are displayed.




Figure 7: The Rule window. Here a metarule of the task Pursue-Hypothesis is displayed.

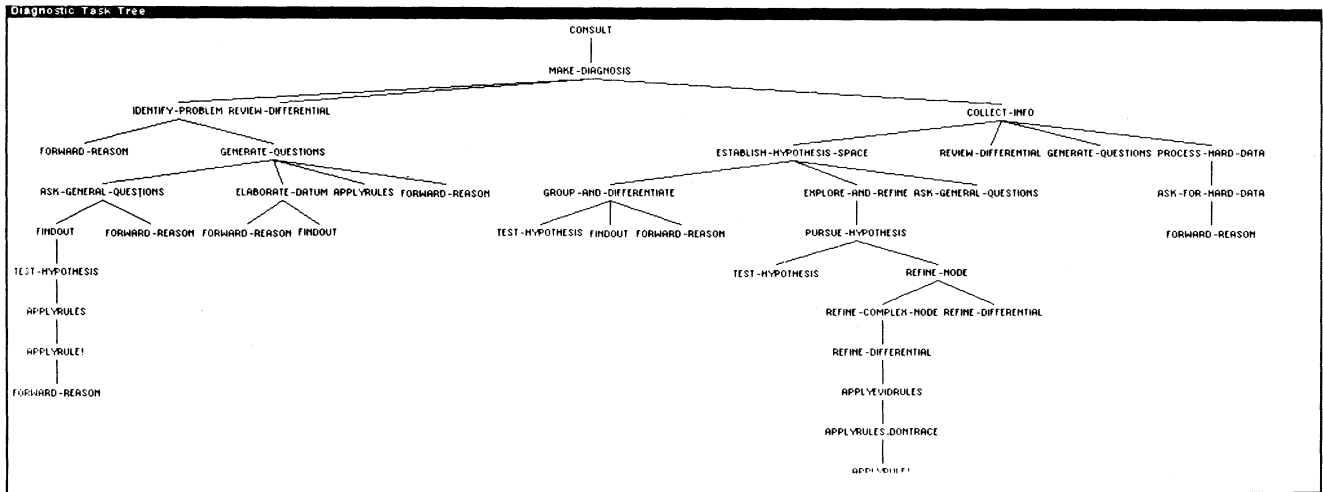


Figure 8: The Diagnostic Task Tree window. When a task is selected in this window, a menu pops up that allows the user to display either the properties of the task in the Task Property window (Figure 5) or the metarules that a task applies in the Metarules window (Figure 6). During a consultation the user can also choose to see dynamic information about task calls. This is described in the section on Dynamic Task windows.

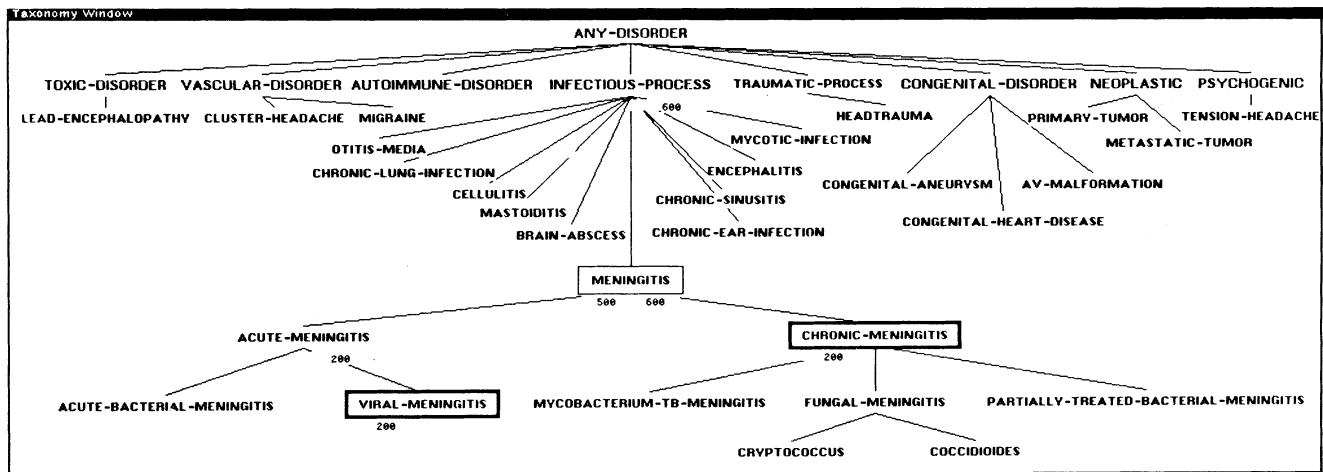


Figure 9. The Taxonomy window. The boxing, flashing, and printing techniques used in this window make the dynamic search strategy visible by displaying dynamic information on top of static knowledge structures, in this case the etiological taxonomy. The differential shown here is NEOMYCIN's internal differential and may not correspond precisely to a physician's differential. The differential shown to students may differ from NEOMYCIN's internal list. This graph was generated, edited, and stored in the same manner as the graph in Figure 3. Note that Figures 9-12 correspond to the same state of the consultation as displayed in Figure 3.

ancestors are also combined; therefore, evidence against infection can decrease the CUMCF of meningitis.)

When the CF or CUMCF is updated for a hypothesis, new values are printed below the node label corresponding to the hypothesis. The CF is printed on the left, the CUMCF on the right if it differs. The figures in this article show the CFs printed as integers from -1000 to +1000; this is how they are represented internally. These numbers are far from precise and should be interpreted as falling into several categories: definite, strongly suggestive, suggestive, weakly suggestive, or no evidence for (or against) a hypothesis. For students the internal CF values will not be printed; instead a graphic notation, such as zero to four pluses or minuses, could be used to indicate the degree of belief.

In the case in which a hypothesis window is not open, the printing, boxing, and flashing of nodes is not done immediately. However, whenever the Taxonomy or Causal Relations window is opened during a consultation, the window is updated so that all the hypotheses are appropriately boxed, and certainty factors are printed. Therefore, the user is free to open these windows at any time. We describe several other windows that display dynamic information.

The Consultation Typescript window (Figure 3). The Consultation Typescript window is opened when a user starts a consultation using the *Consult* menu. This window displays the questions that are asked during a consultation.

Each question is followed by a response that is either supplied by the user or is retrieved automatically from a patient data file. Before an answer is retrieved, the program pauses. The user can then use the mouse to select items or open any windows. A menu is provided that allows a user

to proceed one or more questions further, receive textual explanations, resize the consultation window, and so on.

Evidence for MENINGITIS			
FINDING	RULE(S)	MAXCF	MINCF
TENSE-FONTANEL	RULE060	800	
SEIZURES	RULE060	800	
REDFLAG-CNS-FINDI	RULE323	700	
STIFF-NECK-ON-FLE	RULE424	500	
	RULE183	500	
HEADACHE	RULE424	500	
NEONATE	RULE183	500	
WBC	RULE131	-700	
CSFCCELLCOUNT	RULE131	-700	
	RULE117	-800	
CSFPROTEIN	RULE117	-800	

Figure 10. The Evidence window. The findings and hypotheses displayed in this window are ordered so that the ones that may be most suggestive (have the highest MAXCF) are on top.

RULE424	
If:	1) The patient has a stiff neck on flexion, and
	2) The patient has a headache
Then:	If you are considering infectious-process, there is suggestive evidence (.5) that the patient's infection is meningitis

Figure 11: Here the Rule window displays a domain rule.

Positive Findings		
FINDING	VALUE	CF
AGE	42	
SEX	FEMALE	
RACE	LATINO	
HEADACHE	YES	
STIFF-NECK-ON-FLEXION	YES	
NAUSEA	YES	
HEADACHE-DURATION	10	
HEADACHE-SEVERITY	3	
CNS-FINDING	YES	
STIFF-NECK-SIGNS	YES	
HEADACHE-CHRONICITY	CHRONIC	800
	SUBACUTE	300
CNS-FINDING-DURATION	10	

Figure 12: The Positive Findings window.

Differential		
HYPOTHESIS	CF	CUNCF
VIRAL-MENINGITIS	200	
FUNGAL-MENINGITIS	---	
MYCOBACTERIUM-TB-MENINGITIS	---	
PARTIALLY-TREATED-BACTERIAL-MENINGITIS	---	

Figure 13: The Differential window. This window is displayed several questions after the point shown in Figure 3. Subsequent figures all show windows as displayed at this later point.

Evidence window (Figure 10). This window can be displayed without running a consultation. However, during a consultation dynamic information is overlaid onto static knowledge structures to show the current evidence relations between findings and hypotheses. All potential evidence for a hypothesis is displayed as a table in this window. The first column lists findings and hypotheses that suggest a hypothesis. The second column lists the rules that use these findings or hypotheses to make conclusions about the hypothesis. The third column shows the maximum CF in the rule's action, and the fourth column shows, if different, the minimum CF in the rule's action. Findings, hypotheses, and rules can be selected with the mouse to get more information. For example, a rule's premise and conclusion can be displayed in the Rule window (Figure 11).

During a consultation, GUIDON-WATCH employs boldfaced text and grayed-over regions to provide the user with additional information. The user may have displayed the evidence of meningitis because it was boxed in the Taxonomy window (Figure 9). Seeing that rule 424 succeeded (which is indicated by the bold text), the user can display the rule's premise and conclusion in the Rule window (Figure 11). A finding with a positive value is displayed in bold; a negative finding is grayed over. Analogously, rules that have succeeded are printed in bold; rules that have failed are grayed over. Findings and rules that appear in normal print have not been investigated yet. This simple notation is an effective means of providing a great deal of information in a concise and understandable manner. Furthermore, it illustrates how dynamic information can be displayed on top of static knowledge structures that are displayed in a table format.

Positive Findings window (Figure 12). The Positive Findings window displays all the findings that have a positive value (i.e., the value is "yes," a number, or symbolic). Findings are printed in the first column, values in the second column, and CFs in the third (printed only if less than 1000). Findings are selectable, and when buttoned a pop-up menu is displayed. For example, a user may want to select a finding to see which hypotheses the finding may suggest.

In this window items are printed incrementally during a consultation. If the Positive Findings window is open, new positive findings are printed in the window as soon as they are known. If the window is closed, the whole list of positive findings is printed when the window is opened. This feature provides flexibility for the user, who can open or close the window at any time during a consultation.

Differential windows (Figure 13). Hypotheses on the differential are boxed when they appear in certain windows. However, the differential is such an important structure that a special window is provided for its display (Figure 3).

Hypotheses-with-Evidence windows (Figure 14). However, not all hypotheses for which there is positive evidence are on the differential at a given time. This group includes hypotheses for which there is direct evidence (i.e., at least one rule concluded the hypothesis) and those for which there is belief when propagation is included (i.e., the CUMCF is above a certain threshold). Note that some of these hypotheses may not be on the differential at a given time, and additionally, hypotheses on the differential may not have evidence supporting them.

In both the Differential and Hypotheses-with-Evidence windows, hypotheses that have direct evidence supporting them are printed in bold. These windows also contain columns for CF and CUMCF values. As usual, the hypotheses are selectable. These two windows, as well as the Taxonomy and the Causal Relations windows, illustrate how GUIDON-WATCH provides *multiple views* of the same knowledge structures.

Dynamic Task windows. These windows provide users with dynamic views of the diagnostic strategy as it is instantiated during a consultation. This is a challenging presentation problem because the abstract nature of the diagnostic procedure as it is represented in the task and metarules is not nearly as intuitive to people as are disorder hierarchies, causal networks, domain rules, and lists of findings. Although the goal is to provide a view of NEOMYCIN's reasoning that is understandable to medical students, the model of the diagnostic strategy is in the form of a complex procedure that is intimately tied to basic concepts of computing. For example, task calls are very similar to procedure calls; a task may have a focus and local variables. A focus consists of one or more findings, hypotheses, or rules depending on the task. For example, Test-Hypothesis may have meningitis or another disease as a focus in NEOMYCIN. Tasks invoke other tasks in a chain, similar to procedure calls.

The three windows described here provide a different view of the dynamic diagnostic strategy by using three different graphic formats: a stack, a tree, and a table. Although it has not yet been decided how they will be adapted for instruction, they are already very useful for programmers trying to debug or understand NEOMYCIN's behavior. Programmers can use these windows to find out exactly what NEOMYCIN is doing or has done at a detailed strategic level. Consistent with Model's¹⁴ recommendations, these windows provide monitoring and debugging tools at a level that corresponds to the program's design (e.g., tasks and metarules). This is a great improvement over examining the low-level Lisp stack, which reflects the strategy only in a very indirect way.

Task Stack window (Figure 15). This window displays the current stack of task calls, which is similar to a stack of procedure calls. Its current design shows the tasks in the order that they were called, with the first task printed at the

top of the window. If the task has a focus, it is printed in square brackets after the task. The metarule that the task successfully applied is printed below the task. Metarules are *attached* to the task they invoke by a vertical line. Different font faces are used to distinguish tasks, meta-rules, and foci from one another. Every rule, finding, and hypothesis in the Task Stack window is selectable so that the user can quickly get more detailed information on an item of interest.

Hypotheses With Evidence		
HYPOTHESIS	CF	CUMCF
INFECTIOUS-PROCESS	700	880
MENINGITIS	500	600
CHRONIC-MENINGITIS	200	
VIRAL-MENINGITIS	200	
ACUTE-MENINGITIS	---	200

Figure 14: The Hypotheses with Evidence window.

TaskStack Window	
MAKE-DIAGNOSIS []	
RULE384	
COLLECT-INFO []	
RULE062	
ESTABLISH-HYPOTHESIS-SPACE []	
RULE586	
EXPLORE-AND-REFINE []	
RULE163	
PURSUE-HYPOTHESIS [MYCOBACTERIUM-TB-MENINGITIS]	
RULE171	
TEST-HYPOTHESIS [MYCOBACTERIUM-TB-MENINGITIS]	
RULE603	
APPLYRULES [RULE366 RULE309 RULE309 RULE002 RULE525]	
RULE094	
APPLYRULE! [RULE309]	
RULE095	
FINDOUT [STEROIDS]	
RULE153	
FINDOUT [IMMUNOSUPPRESSED]	
RULE153	
FINDOUT [SYSTEMIC-COMPROMISED]	
RULE153	
FINDOUT [COMPROMISED]	
RULE169	

Figure 15: The Task Stack window. By examining the task stack, the user can see that NEOMYCIN is testing the hypothesis mycobacterium-TB-meningitis. As a result, a rule was applied that led to a series of calls to the task Findout. The last call with the focus "compromised" finally resulted in a question to the user: "Is Mary a compromised-host?" The user would see this in the Consultation Typescript window.

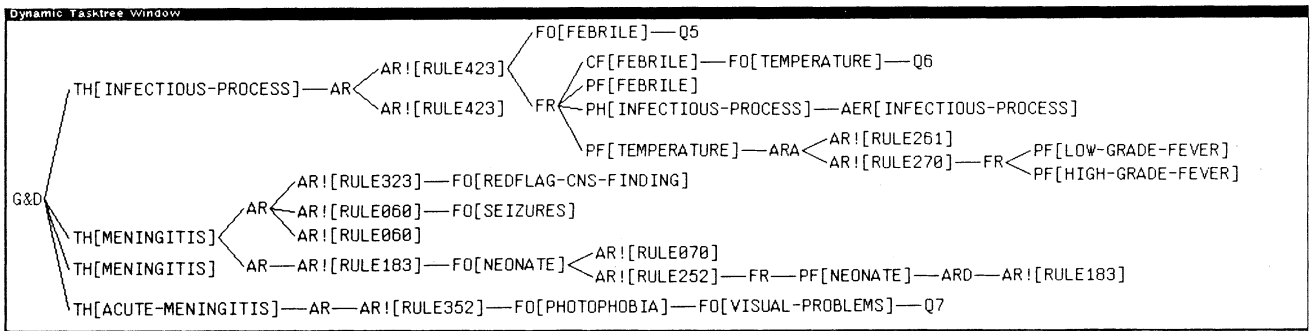


Figure 16: The Dynamic Task Tree window. The node labels of tasks in this tree are abbreviated, but the user can see the full name expanded when a node is selected. TH is short for Test-Hypothesis in this tree. This figure illustrates how the user can see multiple calls of a task and the resulting events in the Dynamic Task Tree window.

The Task Stack window provides a view of the current path through the diagnostic tree with metarules and foci instantiated. By examining the task stack, the user can understand the reason for the current strategy. For example, the user may be interested in why a question is being asked. Students will be able to get textual explanations (Figure 3) that should satisfy their needs, but programmers may want to examine the task stack to understand the computational reasons for a data request at the task and metarule level (see Figure 15).

Dynamic Task Tree window (Figure 16). This window displays a graph that shows all or part of the dynamic history of task calls. This allows a user to view the overall structure of the diagnostic strategy that NEOMYCIN is using during or at the end of a consultation. This capability is useful because the static Diagnostic Task Tree window (Figure 8) shows all possible paths in the task tree; this window shows only the paths that are part of an actual diagnosis and reveals patterns of multiple calls of the same task.

Task History window (Figure 17). This window contains a table of all the invocations of any given task during the consultation. It provides an alternate view (i.e., a slice at a time) of the information displayed in the Dynamic Task Tree window. In the first column the invocation number of the task is printed, with the digit 1 meaning the first time the task was called. In the second column the focus of the task call is printed; in the third column the metarule that

Task History of TEST-HYPOTHESIS			
NO.	FOCUS	CALLER	METARULE
1	INFECTIOUS-PROCESS	G&D	RULE393
2	MENINGITIS	G&D	RULE400
3	MENINGITIS	G&D	RULE400
4	ACUTE-MENINGITIS	G&D	RULE400
5	CHRONIC-MENINGITIS	PUH	RULE171
6	MYCOBACTERIUM-TB-MEN	PUH	RULE171

Figure 17: The Task History window.

invoked the task is printed; and in the fourth column the calling task is printed. As usual, rules, findings, hypotheses, and tasks are selectable. Additionally, the user can select an invocation number to display more information on the history of that task call. For example, the user can display a dynamic task tree with the chosen task invocation as the root.

Together, the three dynamic task windows provide a powerful aid for inspecting the current and past diagnostic strategy used during the consultation. It is clear that medical students would need some instruction in these concepts before the dynamic task windows would be meaningful to them. Part of the problem is that many of the task names are not commonly used in medicine. It is hoped that some of the problem can be alleviated by choosing names for the diagnostic tasks that are more familiar to students. Additionally, some tasks may be hidden from a student's view because they involve computational details of interest to a programmer only.

Prior and related work in graphic interfaces

GUIDON-WATCH was influenced by a diverse collection of work stretching back to Vannevar Bush's seminal article, in which a desk-size, electronic information device called a "memex" was proposed.¹⁵ Doug Engelbart and his colleagues pioneered much of the early work on information-handling systems and provided the basis for the user interfaces commonly found on today's workstations.¹⁶⁻¹⁹ Alan Kay led the Learning Research Group at Xerox PARC that brought similar ideas to fruition on personal workstations.²⁰ Engelbart's group and the LRG shaped a view of the computer as a communications medium by which a user can store, retrieve, manipulate, and transfer information with ease. The LRG's vision of a dynabook²¹⁻²⁴ still remains an exciting dream in the spirit of Bush's memex.

The 1970's also brought many advances in artificial intelligence including the development of knowledge-based systems such as MYCIN. Starting from a different point of

view, Seymour Papert led a group at the Massachusetts Institute of Technology that explored the use of computer languages such as LOGO to teach subjects such as geometry and physics in a new way.²⁵ Papert offers a provocative view of AI and computers in education; he influenced us to consider how we can provide students with conceptual and software tools to explore computational models. John Seely Brown further inspired us to understand the potential of these ideas; his discussion of *reifying the process* of problem solving is particularly relevant to GUIDON-WATCH.²⁶ GUIDON-WATCH makes the abstract diagnostic procedure used in NEOMYCIN more concrete and visible.

Partly because bit-mapped raster displays have only recently been integrated with AI programming environments, little has been written about graphic interfaces in AI programs. Some notable exceptions include Model,¹⁴ AIPS,²⁷ the ONCOCIN project,^{28,29,30} and the STEAMER project.^{31,32} Model demonstrated that graphic displays can facilitate the monitoring and debugging of complex programs (he used MYCIN for an early demonstration of his work). Tsuji and Shortliffe investigated this idea further by implementing several graphic tools for constructing, monitoring, and debugging ONCOCIN's knowledge base and inference procedures. (ONCOCIN is a system that helps physicians administer experimental cancer therapy.) The ONCOCIN group's strong commitment to the use of interactive graphics has resulted in several graphic interfaces, including the *ONCOCIN Interviewer*,³⁰ a program that helps physicians enter patient data.

STEAMER, an instructional program about power-plant operation, uses interactive color graphics in a knowledge-based simulation. It is interesting to compare the use of interactive graphics in STEAMER and GUIDON-WATCH. STEAMER emphasizes the construction of a visible, interactive, and inspectable simulation. It displays the complex physical processes of a steam propulsion plant. NEOMYCIN, on the other hand, is a computational model of diagnostic reasoning. GUIDON-WATCH provides a user with a visible, interactive, and inspectable model of NEOMYCIN's reasoning processes.

Several knowledge-base browsers similar to GUIDON-WATCH were developed more or less concurrently. For example, ART, KEE, SRL+, LOOPS, S.1, and STROBE provide interactive graphics displays that allow programmers to browse class hierarchies and other general data structures.³³⁻³⁶ Sophisticated graphics editors may be provided; for example, STROBE has an excellent knowledge-base editor.³⁷ However, the browsers provided in these systems are very general and are too complex for end users, because they require an understanding of the underlying knowledge-representation framework. On the other hand, KEE and LOOPS do provide support for creating end-user iconic displays. These can be very useful for displaying the state of a complex device.

GUIDON-WATCH differs from other browsing programs because it is tuned to display specific kinds of

knowledge structures (e.g., those found in a HERACLES system). For example, GUIDON-WATCH can display disease taxonomies, causal networks, evidence for a hypothesis, and positive findings in a way that is appropriate for end users such as medical students. Graphics techniques described here illustrate an abstract diagnostic procedure during its actual use. To paraphrase, if a knowledge base is written for HERACLES, an effective user interface is provided automatically.

Future work

A continuing decrease in the price of hardware will provide more opportunities to use higher resolution screens, interactive pictures, color, animation, and interactive video. Certainly, we have only touched the surface in using graphics for viewing a knowledge-based system. Interactive and animated pictures can illustrate facts and processes. However, implementing interactive graphics displays is time-consuming. There is a need for high-level user interface kits that provide most of the common features that developers now have to implement over and over again. It is probable that an object-oriented programming system will be adopted as an extension to the Common Lisp standard.^{38,39} This system could provide the basis for a generic interface shell for Lisp environments. Two examples of interface packages that successfully use the object-oriented approach include MacApp⁴⁰ and EzWin,⁴¹ the latter is written in *flavors*, the object-oriented language within Zetalisp.⁴²

The discussion so far has focused on what interactive graphics can provide for AI systems. However, AI technology can contribute directly to more intelligent graphic interfaces. Current research topics include user modeling, intelligent presentation, and declarative languages for describing graphical interaction. Mackinlay⁴³ is investigating some of these issues. For GUIDON-WATCH, we decided how to present information and hand-coded it. Instead, Mackinlay's program *reasons on its own* about how to present information. For example, it can decide to present data as a bar chart, a pie chart, a plot chart, a table, or a graph. It can also design several alternative sophisticated presentations from simpler ones and then use heuristics to choose one to display to the user.

Another important aspect of Mackinlay's work is that it uses a knowledge-based approach. Therefore, its reasoning is represented in an explicit, declarative language and not in opaque code. The use of a declarative representation results in programs that are easier to understand and modify. We found that the parts of our display code that are trying to be *smart*, such as the management of windows, are poorly represented in Lisp. The code fails to make the underlying reasoning explicit, and it is difficult to modify. Another advantage of using a declarative representation is that it can be used in multiple ways. Mackinlay's current work addresses only the intelligent presentation problem,

but eventually programs may be able to explain why a particular presentation was chosen. The graphics designer using an intelligent computer aid would want a justification for some design decisions. In a teaching program it would be useful if a program could automatically generate questions regarding a presentation on the screen.

The problems involved in developing intelligent interfaces are certainly very difficult, but if they are going to be solved, it seems likely that user-interface behavior must be represented separately in a declarative language or a database. Because the user interface is becoming an increasingly complex and important component of a software system, there are compelling reasons to make a clean separation between the user interface and the rest of a software system.^{27,44,45} There are several reasons to support such highly modular systems:

- they are easier to maintain and debug;
- they can be customized more easily;
- domain independence is possible; and
- an intelligent reasoning component can be interfaced with less difficulty.

In general, programs can be more attuned to individual users. Some users may prefer different configurations of the screen. The size of the fonts chosen in a window may be too small for some users. Optimizing screen space must not interfere with other concerns such as readability of the screen. Future versions of GUIDON-WATCH should allow users to customize the display to their liking while still providing automatic window-management facilities. User models can play a role in smart interfaces that infer a user's preferences. However, a program must have an explicit model of the user to reason about the user's preferences. We believe that a knowledge-based approach (i.e., using declarative representations) is necessary if an intelligent interface must combine general knowledge about presentation with specific knowledge about a user. This is an area for long-term interdisciplinary research in several areas of computer science, psychology, linguistics, communications, education, and graphic design.

Conclusions

GUIDON-WATCH allows a user to view a knowledge-based consultation system in an efficient way. The program demonstrates how multiple windows, menus, and a mouse can be used to achieve this goal. It also demonstrates that stating a diagnostic procedure in a well-structured rule-based language facilitates developing a graphic interface for viewing and inspecting diagnostic problem-solving behavior. The most important principles learned from this effort are as follows:

1. **Providing multiple views of the same knowledge or behavior can help a user understand a complex system.** Tables, trees, pictures, animation, and other graphic formats can offer these different views. The current prototype

of GUIDON-WATCH has made extensive use of trees and tables to display information in multiple, meaningful ways. Hierarchical relationships are naturally represented as trees, and lists of records with several fields are displayed as tables effectively. There are several important events in NEOMYCIN such as changes in the differential, conclusions about findings and hypotheses, and the task calls. Several windows with different formats can provide different views of these events. However, different classes of users may vary with regard to what constitutes an effective user interface.

2. **The use of bold fonts, boxing and graying items, and other graphics techniques can maximize information content and highlight facts and events in a way that is quickly understandable.** The use of these simple techniques in the Taxonomy and Evidence windows illustrates their effectiveness (Figures 9 and 10).

3. **In well-constrained situations it is possible to manage the display and placement of windows automatically.** Screen space is a precious resource, and each window must be designed, sized, and placed to use space efficiently. However, this is a job that can be cumbersome for a user. Additionally, we want to avoid having a user concentrate on the motor activity of using the mouse to move and place windows on the screen. We believe that there is a fundamental difference between the constrained information-retrieval task that GUIDON-WATCH is designed to perform and more creative and open-ended tasks such as programming or writing. For the latter category, the availability of overlapping windows that are usually shaped and positioned under the user's control may be more desirable.

By displaying information in multiple ways and allowing a user to browse the dynamic state of a consultation interactively, we have taken a first step toward reifying the process of reasoning during a NEOMYCIN consultation. Subsequent instructional programs now under development will ask students to explain, debug, and augment their own and program-generated problem-solving behavior. They will use graphic displays like GUIDON-WATCH to compare and contrast alternative solutions to problems. ■

Acknowledgments

Diane Warner Hasling, Julie Thompson Richer, Ted Crovello, and the CG&A reviewers contributed valuable comments to earlier drafts of this article. Curt Kapsner and John Macias are helping us adapt NEOMYCIN and GUIDON-WATCH for use with medical students. Their feedback is helping us determine how GUIDON-WATCH and other programs can best benefit medical students.

This research has been supported in part by ONR and ARI Contract N00014-79C-0302 and more recently by the Josiah Macy Jr. Foundation (grant B852005). Computational resources have been provided by the SUMEX-AIM

facility (NIH grant RR00785). This article appeared as Technical Report KSL-85-20, sponsored by the Knowledge Systems Laboratory at Stanford University.

References


1. W. J. Clancey and R. Letsinger, "NEOMYCIN: Reconfiguring a Rule-based Expert System for Application to Teaching," in *Readings in Medical Artificial Intelligence: The First Decade*, W. J. Clancey and E. H. Shortliffe, eds., Addison-Wesley, Reading, Mass., 1984, pp. 361-381.
2. D. H. H. Ingalls, "Design Principles Behind Smalltalk," *Byte*, Vol. 6, No. 8, Aug. 1981, pp. 286-298.
3. L. Tesler, "The Smalltalk Environment," *Byte*, Vol. 6, No. 8, Aug. 1981, pp. 90-147.
4. S. K. Card and A. Newell, *The Psychology of Human-Computer Interaction*, Lawrence Erlbaum Associates, Hillsdale, NJ, 1983.
5. J. D. Foley, V. L. Wallace, and P. Chan, "The Human Factors of Computer Graphics Interaction Techniques," *IEEE Computer Graphics and Applications*, Vol. 4, No. 11, Nov. 1984, pp. 13-49.
6. D. W. Hasling, W. J. Clancey, and G. Rennels, "Strategic Explanations for a Diagnostic Consultation System," *Int'l J. Man-Machine Studies*, Vol. 20, 1984, pp. 3-19.
7. E. H. Shortliffe, *Computer-Based Medical Consultations: MYCIN*, Elsevier, New York, 1976.
8. B. G. Buchanan and E. H. Shortliffe, *Rule-Based Expert Systems*, Addison-Wesley, Reading, Mass., 1984.
9. V. Van Melle, *System Aids in Constructing Consultation Programs*, UMI Research Press, Ann Arbor, Mich., 1981.
10. W. J. Clancey, "Overview of Guidon," *J. Computer-Based Instruction*, Vol. 10, No. 1 and 2, summer 1983, pp. 8-15. (Also in *The Handbook of Artificial Intelligence*, Vol. 2, A. Barr and E. Feigenbaum, eds. William Kaufmann Inc., Los Altos, Calif., 1982.
11. T. F. Thompson and W. J. Clancey, "The CASTER System: An Experiment in Knowledge Acquisition Within a Generic Expert System Shell," tech. report KSL-85-32, Knowledge Systems Laboratory, Stanford University, Aug. 1985.
12. Xerox Corporation, *Interlisp Reference Manual*, 1983.
13. W. J. Clancey, "Classification Problem Solving," *Proc. Nat'l Conference on AI*, Austin, Tex., Aug. 1984, pp. 49-55.
14. M. Model, "Monitoring System Behavior in a Complex Computation Environment," tech. report STAN-CS-79-701, Computer Science Dept., Stanford University, Jan. 1979.
15. V. Bush, "As We May Think," *Atlantic Monthly*, July 1945, pp. 101-108.
16. D. C. Engelbart, "A Conceptual Framework for the Augmentation of Man's Intellect," in *Vistas in Information Handling*, P. W. Howerton and D. C. Weeks, eds., Spartan Books, Washington, D.C., 1963, pp. 1-29.
17. W. K. English and D. C. Engelbart, "Display-Selection Techniques for Text Manipulation," *IEEE Trans. Human Factors in Electronics*, Vol. HFE-8, No. 1, Mar. 1967, pp. 5-15.
18. D. Engelbart, "Advanced Intellect-Augmentation Techniques," SRI project 7079, final report, July 1970.
19. D. C. Engelbart, "Toward High-Performance Knowledge Workers," *Proc. AFIPS Office Automation Conf.*, Apr. 1982, pp. 279-290.
20. A. Kay, "Microelectronics and the Personal Computer," *Scientific American*, Vol. 237, No. 3, Sept. 1977, pp. 230-244.
21. A. Goldberg, "Educational Uses of a Dynabook," *Computers & Education*, Vol. 3, 1979, pp. 247-266.
22. A. Borning, "Thinglab—A Constraint-Oriented Simulation Laboratory," tech. report SSL-79-3, Xerox Palo Alto Research Center, Palo Alto, Calif., July 1979.
23. S. Weyer and A. Borning, "A Prototype Electronic Encyclopedia," tech. report 84-08-01, Computer Science Dept., University of Washington, Seattle, Aug. 1984.
24. L. Gould and W. Finzer, "Programming by Rehearsal," tech. report SCL-84-1, Xerox Palo Alto Research Center, May 1984.
25. S. Papert, *Mindstorms: Children, Computers, and Powerful Ideas*, Basic Books, Inc., New York, 1980.
26. J. S. Brown, "Process Versus Product—A Perspective on Tools for Communal and Informal Electronic Learning," *Education in the Electronic Age*, July 1983. (Proceedings of a conference sponsored by the Educational Broadcasting Corporation, WNET/Thirteen Learning Lab, pp. 41-58.)
27. F. Zdybel, N. Greenfield, M. Yonke, and J. Gibbons, "An Information Presentation System," *Proc. Seventh Int'l Joint Conf. Artificial Intelligence*, Aug. 1981, pp. 978-984.
28. S. Tsuji and E. Shortliffe, "Graphical Access to the Knowledge Base of a Medical Consultation System," *Proc. AAMSI (American Assoc. for Medical Systems and Informatics) Congress 83*, May 1983, pp. 551-555.
29. S. Tsuji and E. H. Shortliffe, "Graphical Access to Medical Expert Systems: I. Design of a Knowledge Engineer's Interface," tech. report KSL-85-11,

- Knowledge Systems Laboratory, Stanford University, July 1985.
30. C. Lane, J. Differding, and E. Shortliffe, "Graphical Access to Medical Expert Systems: II. Design of an Interface for Physicians," tech. report KSL-85-15, Knowledge Systems Laboratory, Stanford University, July 1985.
 31. J. D. Hollan, E. L. Hutchins, and L. Weitzman, "STEAMER: An Interactive Inspectable Simulation-Based Training System," *The AI Magazine*, Vol. 5, No. 2, 1984, pp. 15-27.
 32. A. Stevens, B. Roberts, and L. Stead, "The Use of a Sophisticated Graphics Interface in Computer-assisted Instruction," *IEEE Computer Graphics and Applications*, Vol. 3, No. 2, Mar./Apr. 1983, pp. 25-31.
 33. M. Stefik, D. G. Bobrow, S. Mittal, and L. Conway, "Knowledge Programming in Loops: Report on an Experimental Course," *The AI Magazine*, Vol. 4, No. 3, fall 1983, pp. 3-13.
 34. J. C. Kunz, T. P. Kehler, and M. D. Williams, "Applications Development Using a Hybrid AI Development System," *AI Magazine*, Vol. 5, No. 3, fall 84, pp. 41-54.
 35. C. Williams, "Software Tool Packages: The Expertise Needed to Build Expert Systems," *Electronic Design*, Vol. 32, No. 16, Aug. 9, 1984, pp. 153-167.
 36. M. H. Richer, "Evaluating the Existing Tools for Developing Knowledge-Based Systems," tech. report KSL-85-19, Knowledge Systems Laboratory, Stanford University, May 1985.
 37. E. Schoen and R. G. Smith, "Impulse, a Display-Oriented Editor for Strobe," *Proc. Nat'l Conf. on AI*, American Assoc. for Artificial Intelligence, Aug. 1983, pp. 356-358.
 38. D. G. Bobrow, K. Kahn, G. Kiczales, L. Masinter, M. Stefik, and F. Zdybel, "COMMONLOOPS—Merging COMMON LISP and Object-Oriented Programming," tech. report ISL-85-8, Xerox Palo Alto Research Center, Palo Alto, Calif., Aug. 1985.
 39. G. L. Steele, *Common LISP—The Language*, Digital Press, Burlington, Mass., 1984.
 40. L. Tesler, ed., *MacApp*, Release 0.1, Apple Computer, Inc., Cupertino, Calif., 1985.
 41. H. Lieberman, "There's More to Menu Systems Than Meets the Screen," *Computer Graphics*, Vol. 19, No. 3, July 1985, pp. 181-189.
 42. D. Weinreb and D. Moon, *Lisp Machine Manual*, Symbolics Inc., Cambridge, Mass., 1981.
 43. J. Mackinlay, "Intelligent Presentation: The Generation Problem for User Interfaces," tech. report HPP-83-34, Computer Science Dept., Stanford University, Mar. 1983.
 44. R. G. Smith, G. M. E. Lafue, E. Schoen, and S. C. Vestal, "Declarative Task Description as a User-Interface Structuring Mechanism," *Computer*, Vol. 17, No. 9, Sept. 1984, pp. 29-38.
 45. E. Ciccarelli, "Presentation-Based User Interfaces," tech. report AI-TR-794, Artificial Intelligence Laboratory, Massachusetts Institute Technology, Aug. 1984.



Mark H. Richer is a scientific programmer III with the Knowledge Systems Laboratory of the Computer Science Department at Stanford University. Previously he coordinated a microcomputer lab at a public elementary school. He has designed and implemented educational software, developed computer curriculum materials, and taught students and teachers how to use computers. His current interests include knowledge-based systems, graphic interfaces, and instructional software systems.

Richer received the MS in computer science and the MA in interactive educational technology, both from Stanford University. He is a member of the IEEE, the ACM, and the AAAI.



William J. Clancey is a senior research associate in computer science at the Knowledge Systems Laboratory of Stanford University. He has been active in expert systems research since he joined the MYCIN project in 1975, for which he was codeveloper of the antibiotic therapy and question-answering programs. His interests lie in computational modeling of problem solving and the design of architectures for expert systems to facilitate construction, explanation, and multiple use. Clancey has published widely on tutoring and expert system methodology and is the coauthor (with E. H. Shortliffe) of *Readings in Medical Artificial Intelligence: The First Decade*.

He received the BA degree in mathematical sciences from Rice University in 1974 and the PhD in computer science from Stanford University in 1979. He is associate editor (Expert Systems) of *IEEE Transactions on Pattern Analysis and Machine Intelligence*. He is also a cofounder and member of the Technical Advisory Board for Teknowledge, Inc.

The authors' address is Stanford University Knowledge Systems Laboratory, Department of Computer Science, 701 Welch Rd., Palo Alto, CA 94304.