An Intelligent Assistant for Patient Health Care

Silvia Miksch, Kenneth Cheng, Barbara Hayes-Roth

Knowledge Systems Laboratory (KSL), M/C 9020, Gates Computer Science Building 2A
Stanford, California 94 305, email: miksch@hpp.stanford.edu

Abstract

The Patient Advocate is designed to be an intelligent assistant for patient-centered health care. Residing on a home computer or special-purpose device and operating within an extended health-care information network, the Patient Advocate will extend medical expertise into the outpatient setting. It will have remote access to the patient's medical record, an understanding of the patient's health status and history, and a model of the patient's interest in health-related issues, preferences for modes and contents of interaction, etc. The Patient Advocate is being designed to provide three kinds of functions. First, it will assist the patient in managing continuing ambulatory conditions, for example chronic problems such as diabetes, special normal conditions such as prenatal care, and wellness issues such as diet, exercise, and stress. Second, it will provide health-related information by allowing the patient to interact with the on-line health-care information network and scan media resources to suggest information of interest. Third, it will act as a remote triage point for clinical services by coordinating patient-relevant information such as reminding the patient when a visit to the clinic is indicated.

We describe a prototype of the Patient Advocate which is designed to support obstetrics patients at risk of gestational diabetes. It is important that the Patient Advocate is platform-independent, runs on widely available hosts, and has access to various Internet resources. Therefore, it is implemented in Java and is Internet accessible.

1. Patient-Centered Care is Needed

In recent years, several knowledge-based systems have been introduced to support health care providers (mainly clinicians) with the monitoring of critical care patients and to assist them with diagnostic decisions and therapy planning (Uckun 1994b). These systems range from simple intelligent alarms to sophisticated systems for anesthesia monitoring or ventilator management (e.g., GUARDIAN (Hayes-Roth, et al. 1989), NEOGANESH/GANESH (Dojat and Sayettat 1995), SIMON (Uckun, Dawant, and Lindstrom 1993), and VIE-VENT (Miksch, et al. 1993)). Automated reactive planners are a special kind of computer-based support which support the process of protocol-based care over significant periods of time. Two approaches can be distinguished. First, the *prescriptive* approach, in which active recommendations are given according to implicit under-lying guidelines,as in GUARDIAN (Hayes-Roth, et al. 1992) for intensive care monitoring, in VIE-VENT (Miksch, et al. 1993; Miksch, et al. 1996) for artificial ventilation management, in ONCOCIN (Tu, et al. 1989) for oncology, in T-HELPER (Musen, et al. 1992) for AIDS management, and in DILEMMA (Herbert, et al. 1995) as a general architecture. Second, the *critiquing* approach, in which the program critiques the physician's actions rather than recommend a complete course of action on its own, such as HyperCritic (Van der Lei and Musen 1991) and ATTENDING (Miller 1984). This approach concentrates on the user's needs and exploits the assumption that the user has considerable domain-specific knowledge (Miller 1986; Silverman 1992).

These systems were built for the convenience of health care providers and ignored the consumers of health care, namely, the patients. Patients are treated only as objects of medical care.

Health care is one of the most information-rich, but cost-intensive, areas. Preventive care is known to be one of the most effective means of decreasing long-term health care expenses. Poor or less-educated patients cannot afford the expense for adequate preventive care, early interventions, or durative clinical observations and interventions. Nevertheless, there are opportunities to improve the quality of health care through more frequent and accurate self-monitoring and consultation with patients, through more effective communication between patients and health care providers, and through more active patient participation.

A patient-centered approach should concentrate on the individual patient's needs and preferences. A patient-centered system has to make patient-specific and context-specific choices about different styles of presentation, has to provide a user interface according to the patient's needs, expertise, and abilities, has to use patient-oriented vocabulary (e.g., different levels of explanation about the past, current, and future states of the patient's health conditions), and has to perform various reasoning processes. These processes range from data acquisition and validation to sophisticated treatment planning and context-sensitive explanations. Additionally, the system should be used at home, but should still connect the patient to several information resources and facilities. In response to these requirements, we have built a patient-oriented system called **Patient Advocate**, which adapts outpatient care delivery to the individual patient's demands, needs, and preferences.

In the following first section we will illustrate the medical scenario and the patient's participation. The second part will describe the general functionality and the overall architecture of the Patient Advocate project. The third section will concentrate on the most important component of the Patient Advocate project, namely, the monitoring and consulting agent. We will describe its guiding principles, its observation and treatment components, and the sharable underlying task-specific ontology to represent time-oriented treatment guidelines. Finally, we will present our current prototype.

2. The Patient-Centered Medical Scenario

When patients with serious medical problems interact with the health care system, they must perform a variety of challenging and consequential information-processing activities over a period of time. These activities support the patient's understanding or her/his evolving medical situation, cooperation with physicians in making decisions about diagnostic and treatment options, and compliance with the agreed-upon treatment. In performing these activities, the patient often requires or benefits from interaction with different health care professionals, other patients, or support

groups. They, in turn, may interact with various health care information systems to advise the patient effectively.

For example, consider a pregnant woman who is diagnosed as having gestational diabetes mellitus. During each phase of the perinatal period, members of a multidisciplinary health care team perform specific activities to ensure that the patient receives optimal care leading to a healthy outcome for the mother and the baby. The health care team — physician, nurse educator, social worker and nutritionist — work closely with the patient throughout the pregnancy, but during most of the observation period the patient has to monitor herself at home without immediate access to any member of the health care team. For example, hyperglycemia (a higher than normal level of blood glucose) may be detected for the second time in the same week after dinner time. The patient could increase the dose of insulin she typically injects before dinner (insulin reduces the level of blood glucose), could reduce carbohydrate intake (e.g., bread) during dinner, or could exercise more after dinner (e.g., running). Therefore, the patient needs to understand the diagnoses and the recommended activities, the alternative procedures for controlling her blood glucose, her weight gain, etc., and their benefits, risks, and side effects to be able to monitor herself at home.

Making informed decisions requires the patient to comprehend complex and often ambiguous factual information and to weigh many uncertain costs, risks, and benefits in the context of her own preferences. Again, she may obtain information and advice from different physicians, nurses, dietitians, and other professionals. Because GDM can be life-threatening for the mother as well as the baby, the pregnant woman performs these activities under great emotional duress. The patient may turn to former GDM patients or support groups to learn about others' experiences and to obtain more general consultation and emotional support. The pregnant woman needs continuing consultations to encourage self-examination and regular health care provider examinations to assist her in interpreting and responding appropriately to any subsequently detected changes.

Although most physicians and other health care providers actively support patient education and patient participation in choosing among diagnostic and treatment alternatives, their ability to provide information and consultation is limited in several important ways. First, health care providers, especially physicians, have limited accessibility. Clinical visits are short and patients know that physicians' time is valuable. Health care experts are not available at all at many arbitrary times — late in the day or in the middle of the night — when patients have important questions to ask or need to struggle with important decisions. Second, some individual health care providers have restricted knowledge of a given patient's disease, diagnostic and treatment options, medical history, recent medical events or issues, and personal preferences. The patient must integrate information and advice about different aspects of her/his health care from several different sources obtained at different times and places.

In addition, different patients have different information needs and preferences. Some patients desire and can comprehend all the relevant information and prefer to exercise considerable autonomy in the course of their own care, while others prefer to receive a more general accounting of their particular situations and to rely upon the advice of their health care providers. Some patients prefer a purely factual presentation of a given body of information, while others prefer to acquire the same information in case histories of patients who have similar conditions. Some patients are quite capable of scheduling their appointments, determining when they need a consultation, reporting relevant symptoms and asking important questions during office visits, and complying with their prescribed therapies, while other patients need constant reminding, prodding, and advising. To supplement limited resources supporting the comprehension, decision-making, and compliance needs of a diverse population of patients, we are designing and building a computer-based Patient Advocate.

2.1. Scenario: Gestational Diabetes Mellitus

The following requirements determined the choice of a test scenario: (1) high-risk pregnancy in an early state, (2) parameters could be measured and monitored at home, (3) monitoring a pregnant woman with the Patient Advocate at home offers an alternative to staying at a hospital.

Interviewing obstetricians, nurse practitioners, nurse-midwives, and previously pregnant women as well as studying obstetrical textbooks resulted in the choice of gestational diabetes mellitus (GDM). Pregnant women who have never had diabetes before but who have a high level of blood glucose (carbohydrate intolerance) during pregnancy are said to have GDM. Type-I is insulin-dependent and Type-II is non-insulin-dependent. GDM affects about 3 percent of all pregnant women, about 100,000 cases in the United States each year (ACOG 1994). An appropriate controlled observation and treatment of GDM may reduce the risks for several adverse perinatal outcomes including excessive fetal growth and birth trauma, fetal death, and neonatal morbidity, including hypoglycemia and hyperbilirubinemia.

One aim of our approach is to obtain real users' needs and preferences. A cooperation with the leading members of the California diabetes and pregnancy program "Sweet Success" at Stanford University Medical Center (CDAPP 1992) accomplished our aim. We participated in medical consulting meetings and acquired the necessary knowledge to structure the knowledge base and the graphical user interface. We developed mock-ups of the Patient Advocate's user interface. Discussing the mock-ups with concerned medical experts and pregnant women clarified the Patient Advocate's input and output and its basic functionality. As a result we developed a platform-independent, widely accessible graphical user interface which can be tailored to the particular patient's needs and preferences (see Section 5).

3. Patient Advocate

In the Patient Advocate project we are designing, building, and demonstrating cooperative technology to support patients' management of their own health-related behavior on a day-to-day basis at home over a period of time. The project facilitates access to different resources and scheduling tools, such as connecting to health care providers, retrieving medical information on the World Wide Web, consulting support groups or other patients, and scheduling an appointment. It utilizes the patient-centered approach mentioned above. Figure 1 summarizes the several resources administered by the Patient Advocate.

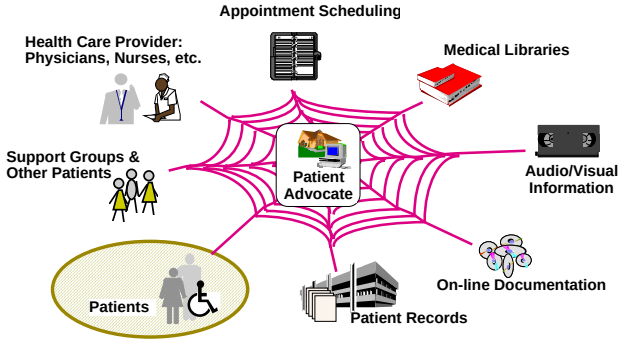

Figure 1: Several resources administered by the Patient Advocate project

3.1. Patient Advocate's Functionality

In general, the Patient Advocate supports the following functionality.

(1) Monitoring and consultation regarding the patient's health condition

The improvement of technical equipment facilitates more frequent and accurate monitoring of patient's health condition. There are several kinds of data about the patient's health condition available from different clinical analyses and devices. All this information is time-stamped and time-varying as a result of changing therapeutic actions and reactions of the human body. The health care provider has to analyze and interpret as much data as possible to derive diagnoses and to recommend appropriate therapeutic actions. In particular the patients have to comply with health care providers' instructions and advice, to exercise health-related common sense, and to monitor themselves for condition-specific danger signs over a period of time. Medical consulting time is expensive. Some patients cannot afford the expense or the time for adequate and regular medical consulting meetings.

Therefore, the central aim of the Patient Advocate project is to assist in clinical practice by helping patients to get a closer insight into their health conditions. Task-adequate methods give a global, comprehensive picture of all information available (e.g., observed parameters available at home, physicians'/nurses' instructions and advice) including explanations about the patients' needs, preferences, and experiences as well as the degree of severity of a situation. Such a comprehensive picture can be achieved by context-sensitive interweaving of different knowledge-based approaches to classify input data and by adequate visualization techniques. Currently, our attention is directed at this interweaving process, even though it is a very complex and partially domain-specific task.

(2) Facilitating access to Web resources

Hundreds of medical information resources around the world are available on the World Wide Web (WWW), including information from federal agencies (e.g., NIH, FDA), clinical guidelines and protocols, literature and library services, medical encyclopedias, continuing medical education resources and, of course, many other non-medical services. Exploring the WWW is a time-consuming and unreliable task because there are too much data available. The variety of candidate sites and paths results in failing to retrieve necessary and accurate information in a reasonable time.

The Patient Advocate will facilitate a context-sensitive access to these various resources, providing the patient with additional explanation or teaching utilities on request. The guiding principle is that only those sites and resources are activated which are meaningful for the current patient's health condition. These sites and resources are annotated with an importance ranking to simplify the search process.

embodies the potential application of particular methods to particular tasks in particular context. Conversely, to function effectively in a particular niche, an agent must exhibit the range of behavior required in that niche and, therefore, must have an architecture that supports the required behavior (Intended Niche => Required Behavior => Sufficient Architecture). Designing the Patient Advocate we started from the intended-use situation (i.e., the intended niche), to support and to consult with patients with respect to their needs and preferences during a period of time; derived the required behavior by participating in medical consulting meetings and interviewing obstetricians, nurse practitioners, nurse-midwives, and previously pregnant women;

(3) Coordinating patient-relevant information

The patient often requires, or benefits from, interaction with different health care professionals, other patients, or support groups. Therefore, the patient is usually confronted with several coordination issues. She/He has to schedule a new appointment with the health care provider, remember scheduled or unscheduled therapy, find other patients with similar clinical conditions or support groups, and more. The Patient Advocate will provide the patients with necessary tools to assist these coordination tasks, such as an appointment scheduler, email connection to the medical staff or other patients, or access to supporting newsgroups.

4. Monitoring and Consulting Agent

In the current phase of the Patient Advocate project we are concentrating on an adaptive agent responsible for monitoring and consulting about the patient's health condition, which represents the most difficult part concerning the reasoning and structuring.

We are using knowledge-based techniques to interpret the raw data, to derive explanations and recommendations. Visualization techniques (Tufte 1983) are applied to display available information according to the patient's needs and preferences. It is important that the Patient Advocate runs on widely available hosts and has access to various Internet resources. Therefore, the Patient Advocate is implemented in Java (Hoff, Shaio, and Starbuck 1996).

4.1. Advice-Giving versus Solution-Presenting

The underlying principle of the monitoring and consulting agent is to guide the system's user toward a *reasonable* explanation or solution, as opposed to presenting a definite solution. In real-life situations people rarely need a solution presented; rather, they need the problems simplified in a way that they can make their own decisions. Experience in providing support to a biotechnical equipment manufacturer showed, that straightforward expert system technology is too narrow (Forslund 1995). The same observations hold in the medical domain. The patients need a global, comprehensive picture of all information available and necessary to judge their health conditions, the recommended activities, and their alternatives. Additionally, the patients require explanations adapted to their preferences, their experiences, and the degree of severity of a situation.

There is a clear relation between the system's intended task, its required behavior, and how it must be constructed. In the domain of adaptive intelligent agents (Hayes-Roth 1995), the agent's architecture determines its potential behavior and, therefore, the niches in which it potentially can function (Agent Architecture => Potential Behavior => Suitable Niches). An "architecture" is the abstract design of a class of agents, a "niche" is a class of operating environments, and a "behavior"

and arrived at a sufficient architecture to support our requirements.

4.2. System Organization

Figure 2 illustrates the agent's system organization and its interactions. The inputs to the monitoring and consulting agent are time-oriented patient data and clinical guidelines. The patient data are acquired from electronic devices or entered by the patient. The clinical guidelines are a course of actions according to a particular clinical condition of a patient, which are entered by the health care providers and have to be represented and transformed in a sharable representation. (The underlying

ontology to deal with the guidelines is explained in the next section.) According to the patient's needs and preferences, particular raw patient data are requested, demanded, and acquired. These raw data are validated to arrive at reliable values. The reliable data are transformed into qualitative descriptions by time-oriented data abstraction methods. The abstracted data are used to interpret the patient's status. The reasoning, based on the sharable clinical guidelines and on the interaction with the patient's needs and tasks, results in context-sensitive recommendations and explanations as well as visualization of the available and necessary data.

4.3. Observation and Treatment Tasks

During the observation and treatment period of GDM the pregnant woman gets instructions from the health care provider to execute several tasks in parallel. Each task takes place at a particular time point and has different frequency and duration (see Figure 3). In the current scenario, the observation period starts after GDM Type-II (non-insulin-dependent) was detected in third trimester pregnancy, as tested by a glucose tolerance test (GTT). In the remaining period of time until delivery, medical screening is executed in conjunction with medication management, nutrition management, observation of the mother's body weight gain, stress management, fetal evaluation, exercise level, and the need for insulin. A suitable component can be added in the case of insulin-dependent (Type-I) diabetes. The various broken lines in Figure 3 illustrate schematically the respective durations and frequencies of each task that should be performed until delivery. During this observation and treatment period the patient has to schedule meetings with the health care providers regularly and to present the self-acquired measurements and her reactions to particular health conditions. The health care providers may add new therapeutic tasks (e.g., injection of insulin) or adapt previous ones. Therefore, both the self-observed data and the applied clinical guidelines change according to the new instructions. (These adaptations are not shown in Figure 3.) The agent is capable of dealing with these new circumstances using a library with alternative clinical guidelines and context-sensitive interpretation techniques.

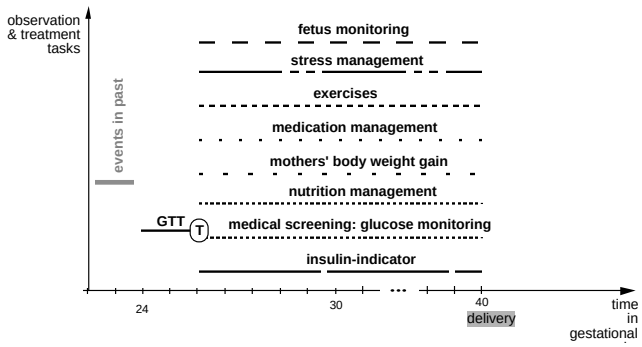

Figure 3: Observation and treatment tasks

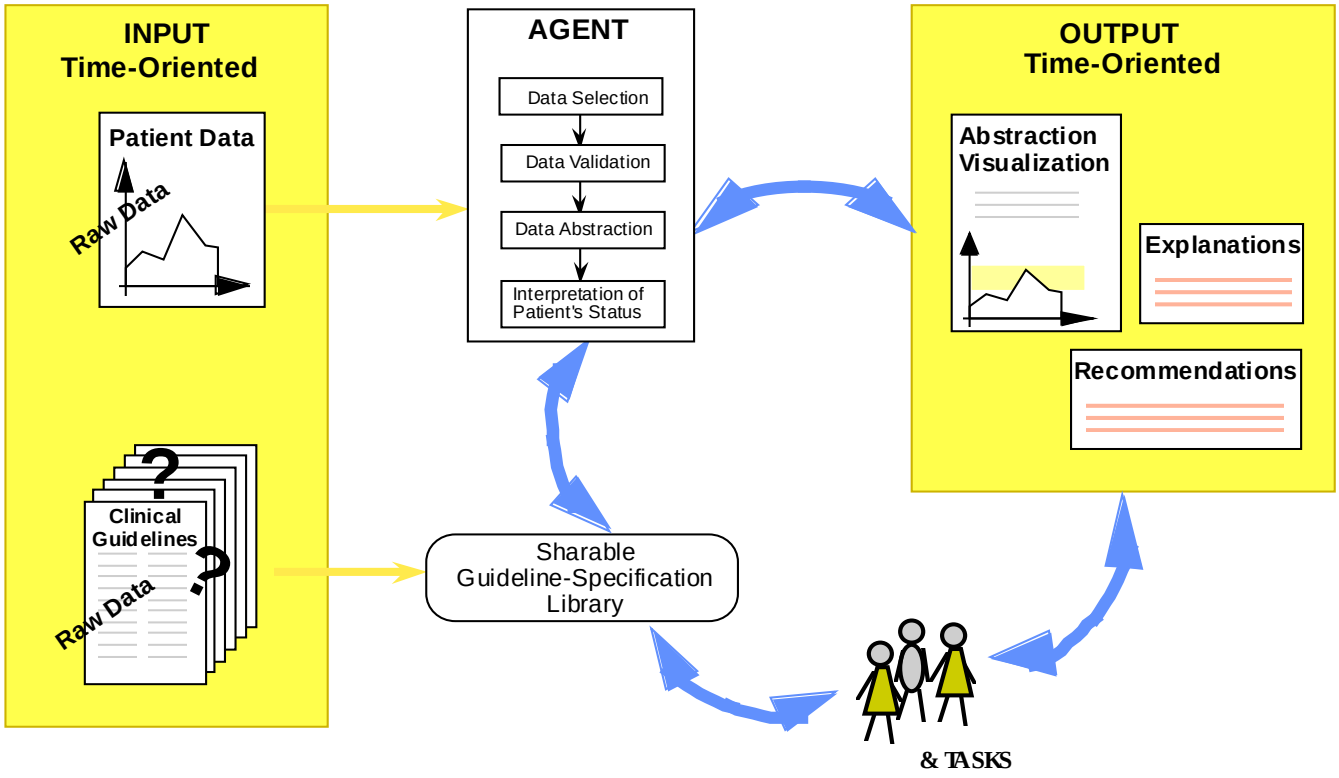

Figure 2: The agent's system organization and its interactions

4.4. Task-Specific Ontology

The underlying knowledge about functional and causal dependencies is vague and ill-structured. Therefore, we are approximating a disease model with clinical guidelines and protocols and customizing it to the individual patient. Clinical guidelines refer to a general principle by which to determine a course of actions according to a particular clinical condition of a patient. Clinical guidelines are often ambiguous or incomplete. For example, a diabetes guideline might recommend a therapy target without any specific recommendations on ways to achieve it. The translation of even well-constructed clinical guidelines of average size into a

knowledge base is a non-trivial task, notable because the context implicit in a guideline must be made explicit for the knowledge structure to be formulated.

In the Asgaard project (Shahar, Miksch, and Johnson 1996b), we have developed a language specific to a set of execution-support tasks (e.g., verification and validation of clinical guidelines, critiquing of an executing agent's action) called Asbru (Shahar, Miksch, and Johnson 1996a). Asbru is an intention-based representation of time-oriented treatment guidelines. The major features of Asbru are that prescribed actions can be continuous; plans might be executed in parallel, in sequence, in a particular order, or periodically; temporal scopes and limit of plans can be flexible; and explicit intentions and preferences can underlie the plan. These features are in contrast to many traditional plan-execution representations (e.g., (Fikes and Nilsson 1971)), which assume instantaneous actions and effects (Uckun 1994a). Finally, Asbru is able to cope with actions and interventions that are often continuous and might have delayed effects.

A clinical guideline represented in Asbru — called a *plan* — consists of a name, a set of arguments, a time annotation, and five components: **preferences**, **intentions**, **conditions**, **effects**, and a **plan body** which describes the actions to be executed. Preferences bias or constrain the selection of a plan to achieve a given goal and express a kind of behavior of the plan (e.g., utility preferences, resource constraints). Intentions are a form of high-level goals and are temporal patterns of actions and external-world states that should be maintained, achieved, or avoided (e.g., intermediate state, overall action pattern). Conditions are temporal patterns that need to hold at particular steps of the plan (e.g., filter condition, suspend condition). Effects describe the functional relationship between the plan arguments and measurable parameters or the overall effect of a plan on parameters. The plan body is a set of plans to be executed in parallel, in sequence, in any order, or with some frequency. All plans are uniformly represented and organized in the *guideline-specification library*. These skeletal plans are a powerful way to reuse existing domain-specific procedural knowledge while leaving room for execution-time flexibility. The plan schemata are instantiated and refined dynamically by the executing agent over significant periods of time and in highly dynamic environments. For a detailed description of the Asgaard project and Asbru see (Shahar, Miksch and Johnson 1996b).

The advantage of the application domain of GDM is that some interesting treatment guidelines are already outlined by the California diabetes and pregnancy program "Sweet Success" (CDAPP 1992). Additionally, cooperation with the leading members of "Sweet Success" program at Stanford University Medical Center supported the acquisition of missing knowledge (e.g., implicit or unmentioned intentions and conditions) to build the guideline-specification library for GDM.

Example: GDM guideline in Asbru

The following represents a part of a guideline used in the CDAPP for controlled observation and treatment of GDM: first, the verbal form (an excerpt from the CDAPP); second, in a graphical version; and third, in the Asbru language.

(1) verbal form

```
Observation and Treatment of Gestational Diabetes
     Mellitus
     (GDM-Type-II)
GLUCOSE MONITORING:
     (after GDM was detected in third trimester
     pregnancy, tested by a glucose tolerance test (GTT)
     being between 140 and 200 mg/dl)
(1)    Patients will check glucose values four times/day
       (fasting and one hour postprandial glucose)
(2)    Preprandial, bedtime and 2 AM blood glucose will
       be added at the discretion of the physician
... (parts deleted) ...
(4)    Treatment goals should be no higher than 130 mg/dl
       for 1-hour post meals, < 100 mg/dl fasting and
       preprandial
... (parts deleted) ...
NUTRITION:
(1)    Patients should be taught a diet based on the
       patients' weight, activity level and number of
       fetuses (regular meals: 3 meals, 3 snacks).
... (parts deleted ) ...
OBSERVE INSULIN INDICATORS
(1) Blood sugar consistently > 100 mg/dl fasting and/or
       one hour postprandial consistently higher than 130
       mg/dl. Attempts at diet modification have failed.
... (parts deleted ) ...
```

(2) graphical version

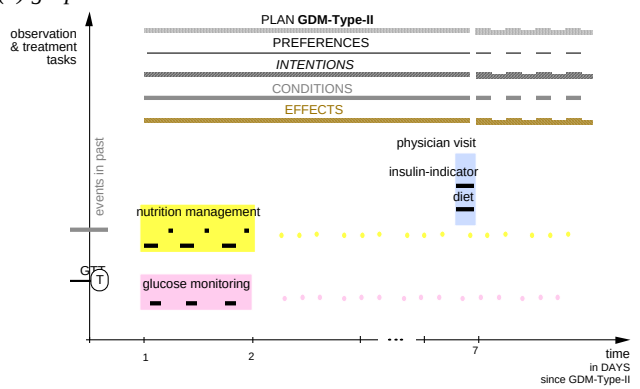

Figure 4: Graphical version of a part of a guideline, used in the CDAPP for controlled observation and treatment of GDM

(3) The Asbru language

The plan body consists of three plans that are executed in parallel. These plans are decomposable into other plans, which exist in the guideline-specification library. Nondecomposable plans are executed by the executing agent. Plan names are written in bold characters. A detailed version of this example in the Asbru language can be found in (Shahar, Miksch and Johnson 1996b).

```
(PLAN observing-GDM-Type-II
      (TIME [[24 WEEKS, 24 WEEKS], [DELIVERY, DELIVERY
],
      [_,_], CONCEPTION])
;; the following time-annotations are local to the GDM
      example
(DOMAIN-DEPENDENT TIME-ASSIGNMENT
   (SHIFTS DELIVERY <- 38 WEEKS)
      ;; time shift from CONCEPTION
   (POINT CONCEPTION
      <-(ask (ARG "what is the conception-date?")))
(ABSTRACTION-ASSIGNMENT
   (CYCLICAL MIDNIGHTS <-[0, 0 HOURS, 24 HOURS]
         BREAKFAST-START-TIME <-[0, 7 HOURS,24 HOURS]))
(PREFERENCES
      (SELECT-METHOD EXACT-FIT)
      (START-CONDITION AUTOMATIC))
(INTENTION:INTERMEDIATE-STATE
   (MAINTAIN blood-glucose-post-meal (<= 130) *
                              ;; raw dara value (no content)
     [[24 WEEKS, 24 WEEKS], [DELIVERY, DELIVERY], [_,_],
        CONCEPTION])
   (MAINTAIN blood-glucose-fasting (<= 100) *
     [[24 WEEKS, 24 WEEKS], [DELIVERY, DELIVERY], [_,_],
        CONCEPTION])
   (MAINTAIN STATE(mothers-body-weight-gain)
     OR SLIGHTLY-LOW NORMAL SLIGHTLY-HIGH) GDM-Type-II
                        ;; a context-specific value
     [[24 WEEKS, 24 WEEKS], [DELIVERY, DELIVERY], [_,_],
        CONCEPTION]))
(INTENTION:INTERMEDIATE-ACTION
```

```
   (MAINTAIN diet regular-meals GDM-Type-II
     [[24 WEEKS, 24 WEEKS], [DELIVERY, DELIVERY], [_,_],
        CONCEPTION]))
(INTENTION:OVERALL-STATE
   (AVOIDED STATE(blood-glucose) HIGH GDM-Type-II
     [[24 WEEKS, 24 WEEKS], [DELIVERY, DELIVERY],
        [7 DAYS,_], CONCEPTION]))
     ;; avoid high blood-glucose level (of any type) for
        more than 7 days
(SETUP-PRECONDITIONS
   (PLAN-STATE one-hour-GTT COMPLETED
     [[24 WEEKS, 24 WEEKS], [26 WEEKS, 26 WEEKS], [_,_],
        CONCEPTION])
(FILTER-PRECONDITIONS
   (one-hour-GTT (140, 200) pregnancy
     [24 WEEKS, 24 WEEKS], [26 WEEKS, 26 WEEKS], [_,_],
        CONCEPTION])
(SUSPEND-CONDITIONS
   (STATE(blood-glucose) HIGH GDM-Type-II
     [[24 WEEKS,24 WEEKS], [DELIVERY, DELIVERY],
        [4 DAYS,_], CONCEPTION]
     (SAMPLING-FREQUENCY 30 MINUTES)))
;; high blood-glucose level (of any type) for at least
4DAYS
(ABORT-CONDITIONS (OR CTIVATED SUSPENDED)
   (insulin-indicator-conditions TRUE GDM-Type-II *
     (SAMPLING-FREQUENCY 30 MINUTES)))
(COMPLETE-CONDITIONS
   (delivery TRUE GDM-Type-II * (SAMPLING-FREQUENCY 30
MINUTES)))
(RESTART-CONDITIONS SUSPENDED
   (STATE(blood-glucose) (OR NORMAL SLIGHTLY-HIGH)
     GDM-Type-II
     [[24 WEEKS,24 WEEKS], [DELIVERY, DELIVERY], [_,_],
        CONCEPTION] (SAMPLING-FREQUENCY 30 MINUTES)))
((PLAN-EFFECTS (GDM-Type-II glucose NORMAL
   ([_, _], [_, _], [20 MINUTES, 60 MINUTES], *NOW*)
0.85))
(DO-ALL-TOGETHER
   (glucose-monitoring)
   (nutrition-management)
   (OBSERVE-insulin-indicators)))
```

4.5. Reasoning Tasks

The monitoring and consulting agent uses the knowledge represented in Asbru for a variety of purposes:

- to execute and recommend a behavior based on a clinical guideline;
- to suggest modifications of guidelines and alternative candidates;
- to observe and determine whether the patient is following the instructions of the health-care provider;
- to explain the patient's health condition on different levels (past and current);
- to project the future development of the patient's health condition.

Therefore, we distinguish between levels of interpretation based on: (1) the **observed data**, (2) the **underlying intentions** of recommended therapeutic actions, and (3) the **expectations** of changing the patient's health condition caused by performed actions. The actual parameter readings (observed data) are used to abstract the patient's health condition. The intentions of therapeutic actions are represented in the clinical guidelines and represent high-level goals. The intentions are used to explain the patient's health condition and are compared with the actual parameter readings. The expectations of parameters are calculated to forecast developments in the patient's health condition. The expectations are represented as two different kinds of effects in the clinical guidelines: first, as functional dependencies between plan arguments and measurable parameters, which are seldom available in the GDM domain; second, as general effects, which describe a qualitative approximation.

5. Prototype: Monitoring and Consulting

Figures 5a and 5b show the parts of the graphical user interface of Patient Advocate's monitoring and consulting agent and illustrate our guiding principles. The interactions with different resources (e.g., connection to the nurse, appointment scheduling, medical libraries) — explained in section 3 — are not displayed in the current prototype. This part is the subject of further research.

We concentrated on an easily understandable interface design minimizing the user's input (e.g., the blood glucose values are automatically entered from the glucometer) and adapting the graphical principles of the CDAPP (e.g., using colors, which are provided in the "Food Guide", to enter the diet plan and the actual intake). To overcome limitations of the screen's size, we grouped parameters context-dependently in separate windows. We incorporated as much knowledge as possible about the interpretation of assessed measurements (e.g., blood glucose, body weight gain) in our visualization of data. Our visualization techniques provided the patients with different views of the observed data to get a closer insight into their clinical condition (e.g., different chart types of plots). Therefore, the pregnant woman can explore her own health condition by obtaining different views of the measurements and getting context-sensitive explanations and recommendation.

Figures 5a and 5b show a particular consulting scenario. Figure 5a displays three observation and treatment tasks: medical screening, nutrition management, and exercise. The lower part shows a menu bar to manipulate the windows. Currently, the medical screening window is activated. Figure 5b shows the corresponding recommendations window. The upper left-hand window in Figure 5a displays the medical screening, currently plotting the blood glucose values measured one hour after dinner for the last 10 days. The patient can choose the different parameters by selecting the "category" button (e.g., fasting glucose or glucose after lunch). The different parameters can be plotted with different chart types (e.g., line graph, histogram) to enable the patient to explore her past and current measurements (scrolling backward and forward is provided by pressing << and >> buttons, respectively). The plot also includes a light gray area which indicates the abnormal region of the measurement. Additionally, the patient can press the help button to get explanations about her health condition. These context-sensitive explanations are generated only by user's requests.

The lower left-hand window shows the nutrition management for the last five days. The nutritionist develops a diet management plan based on the patient's weight, activity level, and number of fetuses, including the recommended nutrient intake for a day. Usually, the diet should include regular meals (three main meals and three snacks) with sufficient nutrients to support the target weight gain. The recommended diet plan is plotted in different gray rectangles using the most important nutrients, namely carbohydrates, protein, and fat. The original interface is plotted in color, therefore the different recommended nutrients are plotted in different colors also. The patient can choose additional parameters, such as vitamins. The thick black lines display the actually-consumed nutrients during the days. The white rectangles estimate the amount which has been burned by exercise. These estimations are approximations, because exact functional dependencies don't exist. Therefore, the patient gets a comprehensive picture of how she had obeyed the medical advice. Again the patient can get explanations about the circumstance by pressing the help button, or choose a different chart type to explore pervious days.

The upper right-hand window displays the exercise for the last ten days. On the vertical axis 24 hours and the corresponding duration of exercise are plotted. We distinguish three levels of activity, namely light, moderate, and heavy. Again, the patient

can scroll to previous days or ask for an explanation by pressing the context-sensitive help button.
The current scenario example shows a pregnant woman with a higher than normal level of blood glucose detected for the second time in the same week after dinner. The patient had eaten too many carbohydrates, too much fat, and too few proteins during some days. She had performed too little exercise to stay within the tolerance level. The recommendation on June 28, 1996 (the last day plotted in the interface) is to change her eating and exercising habits. She should reduce her carbohydrate intake (e.g., bread) during dinner or exercise more after dinner (e.g., running). Additional information and explanations can be demanded by pressing the "detail" button.

6. Evaluation Environments

In the next step we are evaluating the monitoring and consulting agent in three environments. First, the cooperating medical members from the CDAPP at Stanford University Medical will evaluate the soundness of the agent. Second, pregnant women, who have GDM and are involved in the CDAPP, will test the usefulness of the agent. After a revision period in response to the previous two evaluation steps, we will announce the agent on appropriate newsgroups to encourage different kinds of people to evaluate the GDM scenario on the Internet. Currently, we are preparing a test server including CGI scripts, so that a person can log in, store their entered and interpreted data in an user-specific directory, and repeatedly log in to the agent until their account has expired. This procedure enables us to evaluate the agent during a longer time period of observation and treatment, rather than in a single session.

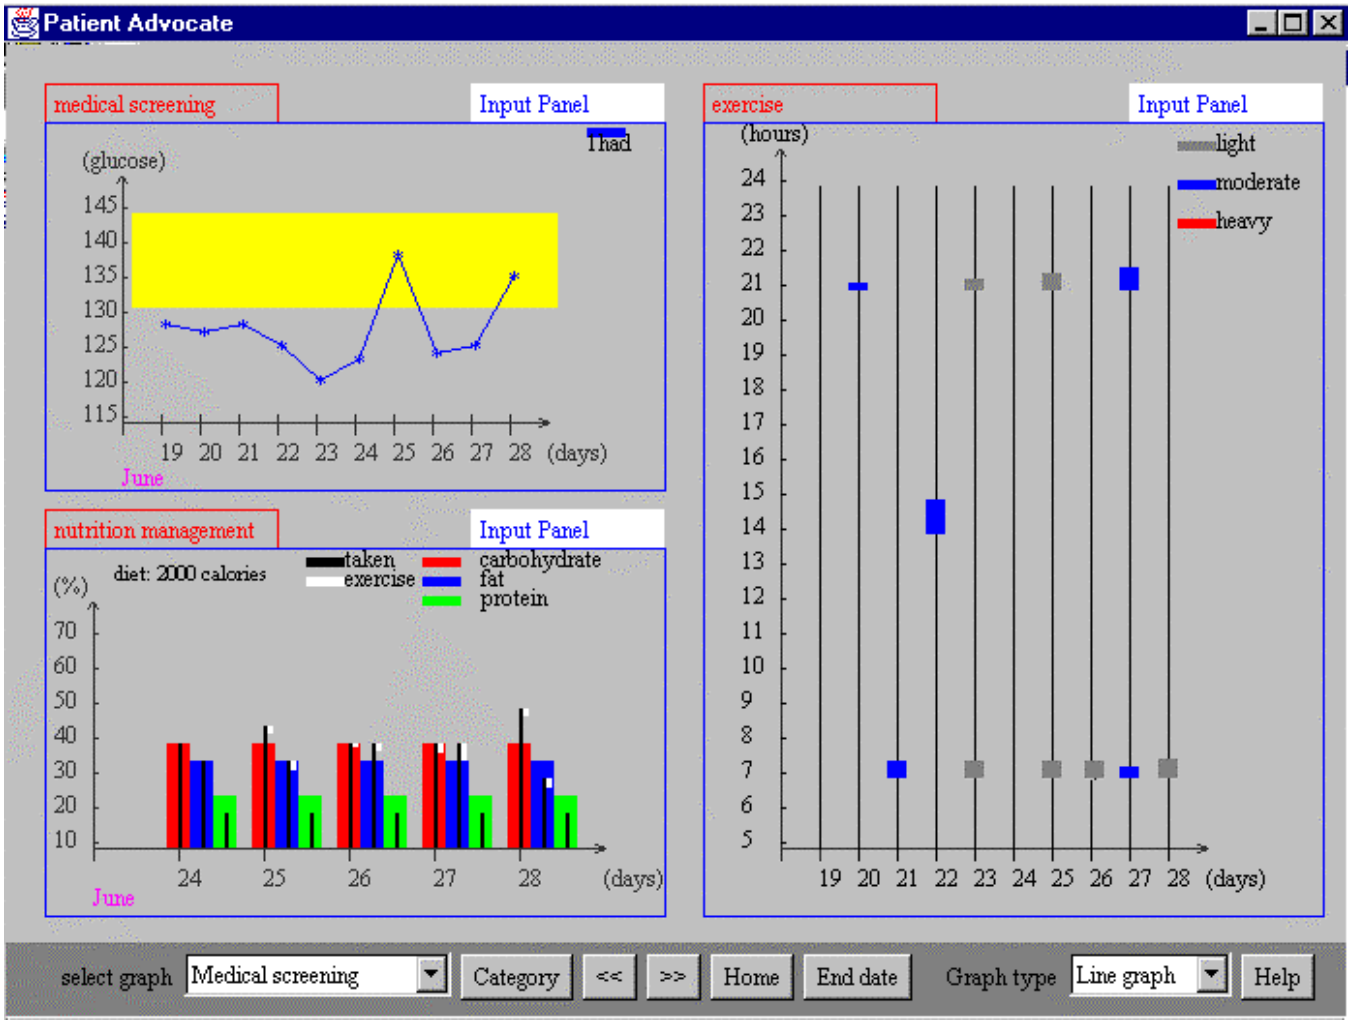

Figure 5a: Patient Advocate's monitoring and consulting agent: medical screening, nutrition management, and exercise screen (scenario GDM Type II)

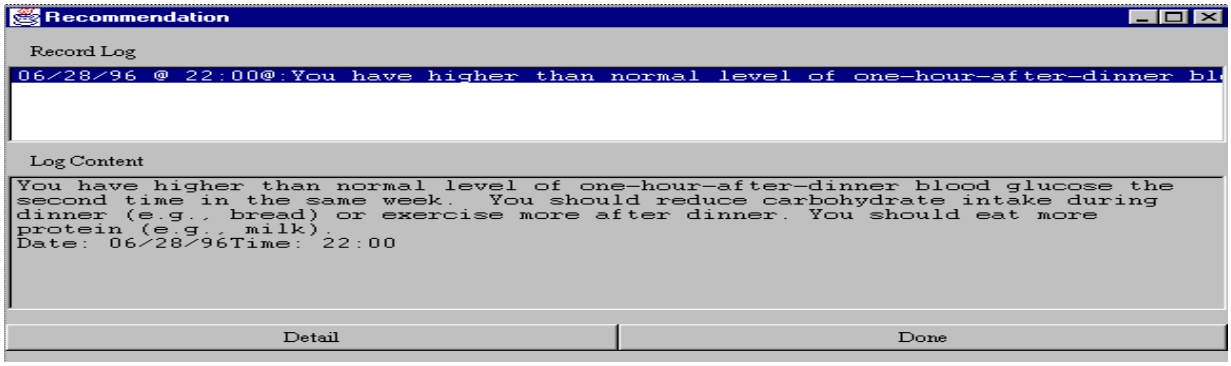

Figure 5b: Monitoring and consulting agent's recommendation panel (scenario GDM Type II)

7. Conclusion

We have demonstrated a prototype of an intelligent assistant for patient-centered health care called Patient Advocate which helps patients to get a clearer insight into their health conditions and to cope with medical instructions and advice on a day-to-day basis at home. It applies knowledge-based, cooperative technologies to interpret raw data, to derive context-sensitive recommendations and explanations, and to visualize the necessary information with different views, depending on the patient's needs and preferences. It is important that the Patient Advocate runs on widely available hosts and has access to various Internet resources. Therefore, it is implemented in Java and is Internet accessible.

Acknowledgments: Patient Advocate is part of the HIIP project (Health Information Infrastructure Program) supported by DARPA Grant N66001-94-D-6055. Silvia Miksch is supported by "Erwin Schrödinger Ausland-stipendium, Fonds zur Förderung der wissenschaftlichen Forschung", J01042-MAT. The authors thank Anne Regenstein, Gretchen D. Flanagan, and Lylia Needham (California diabetes and pregnancy program "Sweet Success") for their fruitful cooperation and medical contributions and Lee S. Brownston for useful comments.

References

ACOG 1994. Diabetes and Pregnancy. *acog Technical Bulletin, An Educational Aid to Obstetrician-Gynecologists*, 200(December).

CDAPP 1992. *SWEET SUCCESS: California Diabetes & Pregnancy Program, Guidelines for Care.*

Dojat, M., and Sayettat, C. 1995. A Realistic Model for Temporal Reasoning in Real-Time Patient Monitoring. *Applied Artificial Intelligence*, 10(2):121-43.

Fikes, R. E., and Nilsson, N. J. 1971. A New Approach to the Application of Theorem Proving to Problem Solving. *Artificial Intelligence*, 2(3/4):189-208.

Forslund, G. 1995. Toward Cooperative Advice-Giving Systems: A Case Study in Knowledge-Based Decision Support. *IEEE Expert*, August:56-62.

Hayes-Roth, B. 1995. An Architecture for Adaptive Intelligent Systems. *Artificial Intelligence, Special Issue on Agents and Interactivity*, 72(1-2):329-65.

Hayes-Roth, B., Washington, R., Ash, D., Hewett, R., Collinot, A., Vina, A., and Seiver, A. 1992. GUARDIAN: A Prototype Intelligent Agent for Intensive-Care Moni-toring. *Artificial Intelligence in Medicine*, 4(2):165-85.

Hayes-Roth, B., Washington, R., Hewett, R., and Seiver, A. 1989. Intelligent Real-Time Monitoring and Control. In Sridharan, N. S. (ed.) *Proceedings of the 11th International Joint Conference on Artificial Intelligence (IJCAI-89)*, Detroit, Michigan, 243-9. Morgan Kaufmann, San Mateo.

Herbert, S. I., Gordon, C. J., Jackson-Smale, A., and Renaud Salis, J.-L. 1995. Protocols for Clinical Care. *Computer Methods and Programs in Biomedicine*, 48:21-6.

Hoff, A. v., Shaio, S., and Starbuck, O. 1996. *Hooked on JAVA* Reading, MA: Addision-Wesley.

Miksch, S., Horn, W., Popow, C., and Paky, F. 1993. VIE-VENT: Knowledge-Based Monitoring and Therapy Planning of the Artificial Ventilation of Newborn Infants. In Andreassen, S., et al. (eds.), *Proceedings of the Artificial Intelligence in Medicine, 4th Conference on Artificial Intelligence in Medicine Europe (AIME-93)*, 218-29. IOS Press, Amsterdam.

Miksch, S., Horn, W., Popow, C., and Paky, F. 1996. Utilizing Temporal Data Abstraction for Data Validation and Therapy Planning for Artificially Ventilated Newborn Infants. *Artificial Intelligence in Medicine*, 8(6):

Miller, P. L. 1984. *A Critiquing Approach to Expert Computer Advice: ATTENDING* Boston: Pitman.

Miller, P. L. 1986. *Expert Critiquing System: Practice-Based Medical Consultation by Computer* New York, NY: Springer.

Musen, M. A., Carlson, C. W., Fagan, L. M., Deresinski, S. C., and Shortliffe, E. H. 1992. T-HELPER: Automated Support for Community-Based Clinical Research. In Frisse, M. E. (ed.) *Proceedings of the Sixteenth Annual Symposium on Computer Applications in Medical Care (SCAMC-92)*, 719-23. McGraw Hill, New York, NY.

Shahar, Y., Miksch, S., and Johnson, P. 1996a. *An Intention-Based Language for Sharing Clinical Guidelines*, Knowledge Systems Laboratory, Stanford University, Stanford, CA 94305, USA, KSL-96-15.

Shahar, Y., Miksch, S., and Johnson, P. 1996b. *A Task-Specific Ontology for Design and Execution of Time-Oriented Skeletal Plans*, Section on Medical Informatics, Stanford University, Stanford, CA 94305, USA, Report SMI-96-0633.

Silverman, B. G. 1992. Survey of Expert Critiquing Systems: Practical and Theoretical Frontiers. *Communication of the ACM*, 35(4):106-27.

Tu, S. W., Kahn, M. G., Musen, M. A., Ferguson, J. C., Shortliffe, E. H., and Fagan, L. M. 1989. Episodic Skeletal-Plan Refinement on Temporal Data. *Communications of ACM*, 32:1439-55.

Tufte, E. R. 1983. *The Visual Display of Quantitative Information* Cheshire, CT: Graphics Press.

Uckun, S. 1994a. *Instantiating and Monitoring Skeletal Treatment Plans*, Knowledge Systems Laboratory, Stanford University, Stanford, CA 94305, USA, KSL 94-49.

Uckun, S. 1994b. Intelligent Systems in Patient Monitoring and Therapy Management: A Survey of Research Projects. *International Journal of Clinical Monitoring and Computing*, 11:241-53.

Uckun, S., Dawant, B. M., and Lindstrom, D. P. 1993. Model-based Diagnosis in Intensive Care Monitoring: the YAQ Approach. *Artificial Intelligence in Medicine*, 5(1):31-48.

Van der Lei, J., and Musen, M. A. 1991. A Model for Critiquing based on Automated Medical Records. *Computers and Biomedical Research*, 24(4):344-78.