Knowledge Acquisition for Clinical-Trial Selection

Savvas Nikiforou, Eugene Fink, Lawrence O. Hall, Dmitry B. Goldgof, and Jeffrey P. Krischer

nikiforo@csee.usf.edu, eugene@csee.usf.edu, hall@csee.usf.edu,
goldgof@csee.usf.edu, jpkrischer@moffitt.usf.edu

Computer Science and Engineering, University of South Florida, Tampa, Florida 33620

Abstract— **When medical researchers test a new treatment procedure, they recruit patients with appropriate medical histories. An experiment with a new procedure is called a clinical trial. The selection of patients for clinical trials has traditionally been a labor-intensive task, which involves the matching of medical records with a list of eligibility criteria, and studies have shown that clinicians can miss up to 60% of the eligible patients. A recent project at the University of South Florida has been aimed at the automation of this task. We have developed an intelligent agent that selects trials for eligible patients. We report the work on the representation and entry of the related knowledge about clinical trials. We describe the structure of the agent's knowledge base and the interface for adding new trials.**

Keywords— **Knowledge representation, medical expert systems, user interfaces.**

I. Introduction

Cancer causes 550,000 deaths in the United States every year, and the treatment of cancer is an active research area. Medical experts explore new treatment methods, such as drugs, surgery techniques, and radiation therapies. An experiment with a new treatment procedure is called a *clinical trial*. When researchers conduct a trial, they recruit patients with an appropriate cancer type and medical history. The selection of patients has traditionally been a manual procedure, and studies have shown that clinicians can miss up to 60% of the eligible patients [12, 22, 30].

A recent project at the University of South Florida has been aimed at automatic selection of patients for clinical trials. We have developed an intelligent agent that prompts a clinician for a patient's data and identifies all matching trials [1, 11]. It includes a knowledge base with information about available clinical trials, criteria for selecting patients, and related medical tests.

We report the work on a web-based interface that enables a clinician to enter new trials without the help of a programmer. We have used the interface to build a knowledge base for clinical trials at the Moffitt Cancer Center, located at the University of South Florida. We review the previous work on medical expert systems (Section II), explain the knowledge representation in the developed agent (Section III), and describe the interface for adding new knowledge (Section IV).

II. Previous Work

Researchers began to work on medical applications of artificial intelligence in the early seventies. Shortliffe and his colleagues developed the MYCIN system, which diagnosed bacterial diseases [5, 25, 26]. Experiments showed the effectiveness of MYCIN, which led to the development of other medical systems [5, 14], such as NEOMYCIN, PUFF, CENTAUR, and VM.

Musen *et al.* built a rule-based system, called EON, that selected AIDS patents for clinical trials [17]. Ohno-Machado *et al.* developed the AIDS^2 system, which also assigned AIDS patients to clinical trials [19]. Bouaud *et al.* created a cancer expert system, called ONCODOC, that suggested alternative trials for each patient and allowed a physician to choose among them [3, 4]. Séroussi used ONCODOC to select participants for clinical trials at two hospitals, which helped to increase the number of selected patients by a factor of three [23, 24].

Early expert systems did not have knowledge-acquisition tools, and programmers hand-coded the related rules. To simplify knowledge entry, researchers implemented specialized tools for some systems [13, 15].

Eriksson pointed out the need for tools that would allow efficient knowledge acquisition, and described a system for building such tools [6]. Tallis *et al.* developed a library of scripts for modifying knowledge bases, which helped to enforce the consistency of the modified knowledge [7, 27, 28, 29]. Kim and Gil considered the use of scripts for building new knowledge-acquisition tools, and created a system for evaluating these tools [9, 10]. Blythe *et al.* designed a general knowledge-acquisition interface based on previous techniques [2].

Musen developed the PROTÉGÉ environment for creating knowledge-acquisition tools [14, 16], which proved effective for the development of knowledge systems, including the AIDS expert systems [20], asthma treatment selection [8], and elevator-design rules [21].

III. Knowledge Base

Physicians at the Moffitt Cancer Center have about 150 clinical trials available for cancer patients. They have identified criteria that determine a patient's eligibility for each trial, and they use these criteria to select trials for eligible patients. Traditionally, physicians have selected trials by a manual analysis of patients' data. The review of resulting selections has shown that they usually do not check all clinical trials and occasionally miss an appropriate trial.

To address this problem, we have built an intelligent agent that helps to select trials for each patient. It prompts a clinician to enter the results of medical tests, and uses them to identify appropriate trials.

In Figure 1(a), we give a simplified example of eligibility criteria for a clinical trial. This trial is for young and

(a) Eligibility criteria

1. The patient is female.
2. She is at most forty-five years old.
3. Her cancer stage is II or III.
4. Her cancer is not invasive.
5. At most three lymph nodes have tumor cells.
6. Either
   - the patient has no cardiac arrhythmias, or
   - all tumors are smaller than 2.5 centimeters.

(b) Tests and questions

General information
What is the patient's sex?
What is the patient's age?

Mammogram, Cost is $150
What is the cancer stage?
Does the patient have invasive cancer?

Biopsy, Cost is $300
What is the cancer stage?
How many lymph nodes have tumor cells?
What is the greatest tumor diameter?

Electrocardiogram, Cost is $200
Does the patient have cardiac arrhythmias?

(c) Eligibility expression

$sex = \text{FEMALE}$ and
$age \leq 45$ and
$cancer\text{-}stage \in \{\text{II}, \text{III}\}$ and
$invasive\text{-}cancer = \text{NO}$ and
$lymph\text{-}nodes \leq 3$ and
$(arrhythmias = \text{NO}$ or
$tumor\text{-}diameter \leq 2.5)$

Fig. 1. Example of eligibility criteria, tests, and questions.

middle-aged women with a noninvasive cancer at stage II or III. When testing a patient's eligibility, a clinician has to order three medical tests (Figure 1b). The agent first prompts the clinician to enter the patient's sex and age. If the patient satisfies the corresponding conditions, the agent asks for the mammogram results and verifies Conditions 3 and 4; then, it requests the biopsy and electrocardiogram data.

The agent's knowledge base includes questions, tests, and logical expressions that represent eligibility for each trial. We give an example of tests and questions in Figure 1(b), and a logical expression in Figure 1(c).

The agent supports three types of questions; the first type takes a yes/no response, the second is multiple choice, and the third requires a numeric answer. For example, the cancer stage is a multiple-choice question, and the tumor diameter is a numeric question. The description of a medical test includes the test name, dollar cost, and list of questions that can be answered based on the test results. For instance, the mammogram in Figure 1 has a cost of $150, and it allows the answering of two questions. Different tests may answer the same question; for example, both mammogram and biopsy show the cancer stage.

We encode the eligibility for a clinical trial by a logical expression, which may include variables that represent the available medical data, as well as equalities, inequalities, "set-element" relations, conjunctions, and disjunctions. For example, we encode the criteria in Figure 1(a) by the expression in Figure 1(c).

The agent collects data until it can determine whether the eligibility expression is TRUE or FALSE. For instance, if a patient's sex is MALE, then the expression in Figure 1(c) is FALSE, and the agent immediately rejects this trial. If the sex is FEMALE, the agent has to ask more questions. If the knowledge base includes many clinical trials, the agent checks a patient's eligibility for each of them. It first asks for the tests related to multiple trials, and then requests additional tests for specific trials.

IV. Entering Eligibility Criteria

We have designed a web-based interface for adding new clinical trials [18], which consists of two main parts; the first part is for adding information about medical tests (Figure 2), and the second is for eligibility criteria (Figure 3). The interface includes ten screens; two of them are "start screens," which can be reached from any other screen. We give an example of entering eligibility criteria, describe the two parts of the interface, and present experiments on its effectiveness.

Example: Suppose that a user needs to enter the criteria shown in Figure 1. First, she utilizes the *"Adding tests"* screen to enter the three tests (Figure 4). Then, she adds the related questions; to enter questions for a specific test, she selects the test and clicks *"Modify"* (Figure 4), and the agent displays the *"Modifying a test"* screen (Figure 5). To add a question, she clicks the appropriate button at the bottom (Figure 5) and then types the question (Figure 6).

After adding the questions for all tests, the user goes to the *"Adding clinical trials"* screen and initializes a new trial (Figure 7). She gets the *"Selecting tests"* screen and chooses the tests related to the current trial (Figure 8). Then, she marks relevant questions and the answers that make a patient eligible (Figure 9). If the eligibility criteria include disjunctions, she has to use the screen for composing logical expressions (Figure 10).

Tests and questions: The interface for adding tests and questions includes six screens (Figure 2). The start screen is for viewing the available tests and defining new ones, whereas the other screens are for modifying tests and adding questions.

We show the start screen in Figure 4; its left-hand side allows viewing questions and going to a modification screen. If the user selects a test and clicks *"View,"* the agent shows the questions related to this test. If the user clicks *"Modify,"* it displays the *"Modifying a test"* screen (Figure 5). The right-hand side of the start screen allows adding a new test by specifying its name and cost.

The *"Modifying a test"* screen shows the information about a specific test, which includes the test name, cost, and related questions. The user can change the test name and cost; the four bottom buttons allow moving to the screens for adding and deleting questions.

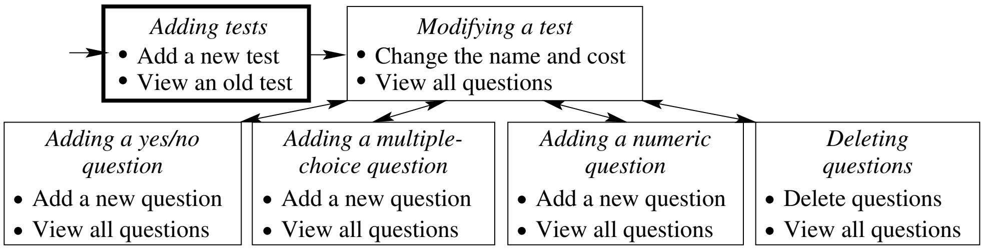

Fig. 2. Entering tests and questions. We show the screens by rectangles and the transitions between them by arrows. The bold rectangle is the start screen.

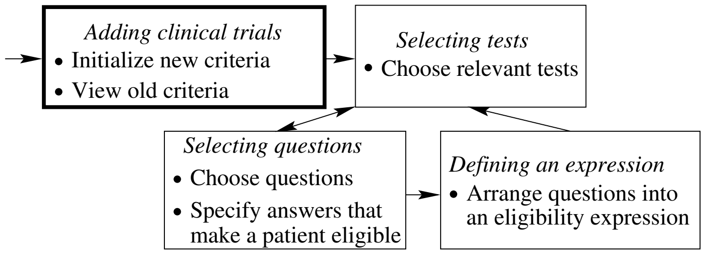

Fig. 3. Entering eligibility criteria.

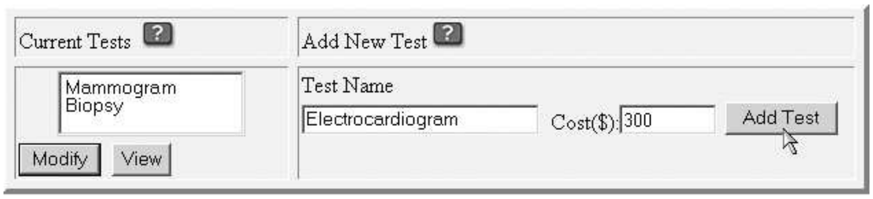

Fig. 4. Adding a new test.

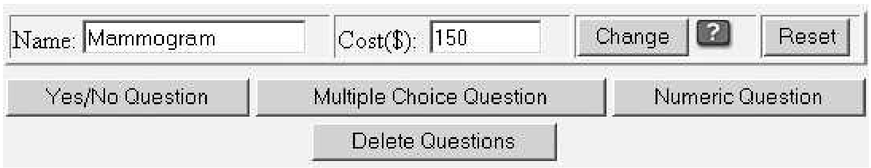

Fig. 5. Modifying a test; the bottom buttons are for moving to question-entry screens.

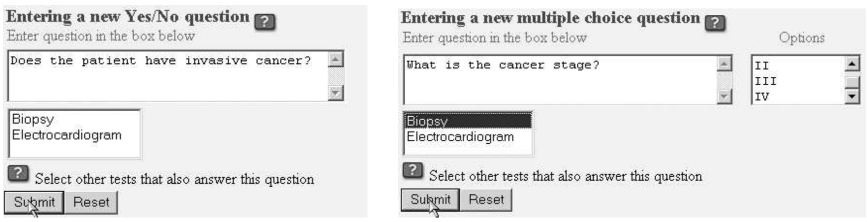

(a) Yes/no question.

(b) Multiple-choice question.

Fig. 6. Adding new questions; the user enters a question and answer options.

Fig. 7. Adding a new clinical trial.

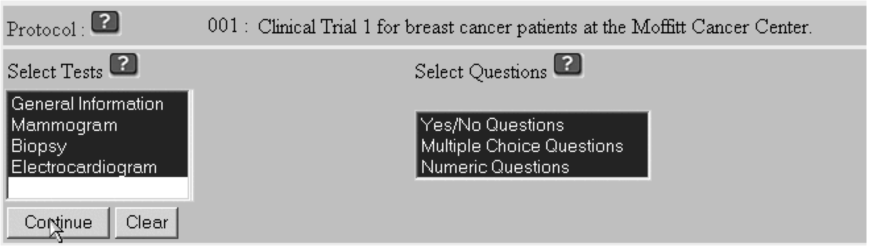

Fig. 8. Choosing tests and question types.

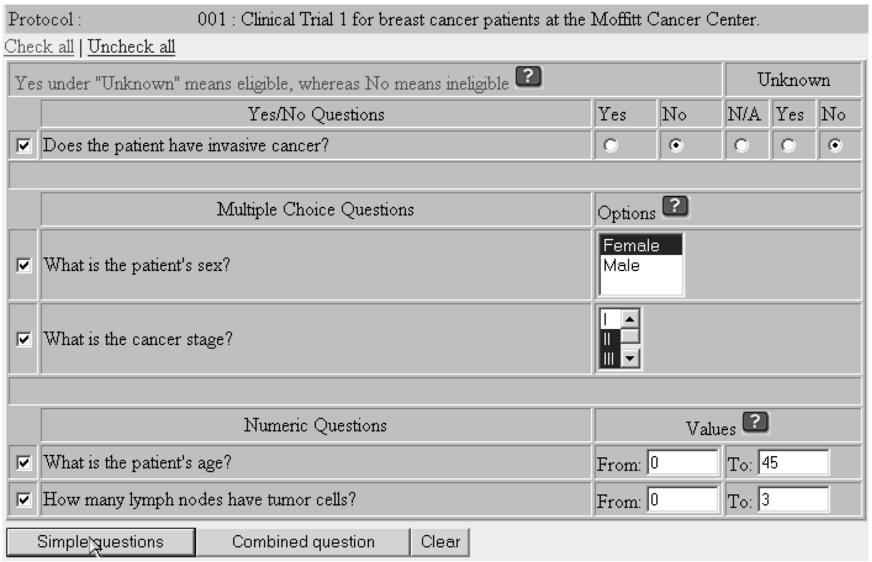

Fig. 9. Selecting questions and answers. The user checks the questions for the current clinical trial and marks the answers that satisfy the eligibility criteria.

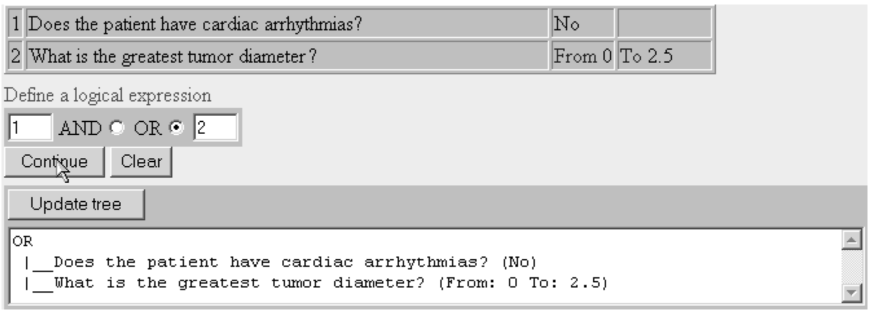

Fig. 10. Combining questions into a logical expression.

We show the screens for adding yes/no and multiple-choice questions in Figure 6; the screen for numeric questions is similar. The user can enter a new question for the current test, along with a set of allowed answers. If the question is also related to other tests, the user has to mark them in the lower box. The *"Deleting questions"* screen is for removing old questions.

Eligibility conditions: The mechanism for entering eligibility criteria consists of four screens (Figure 3). The start screen allows the user to initialize a new clinical trial and view the criteria for old trials. If the user needs to modify a clinical trial, the agent first displays the test-selection screen (Figure 8). The user then chooses related tests and question types, and clicks *"Continue"* to get the question list.

The next screen (Figure 9) allows the user to select specific questions and mark the answers that make a patient eligible. For a multiple-choice question, the user may specify several eligibility options; for example, a patient may be eligible if her cancer stage is II or III. For a numeric question, the user has to specify a range of values; for instance, a patient may be eligible if her age is between 0 and 45 years. If the user clicks *"Simple questions,"* the agent generates a conjunction of the

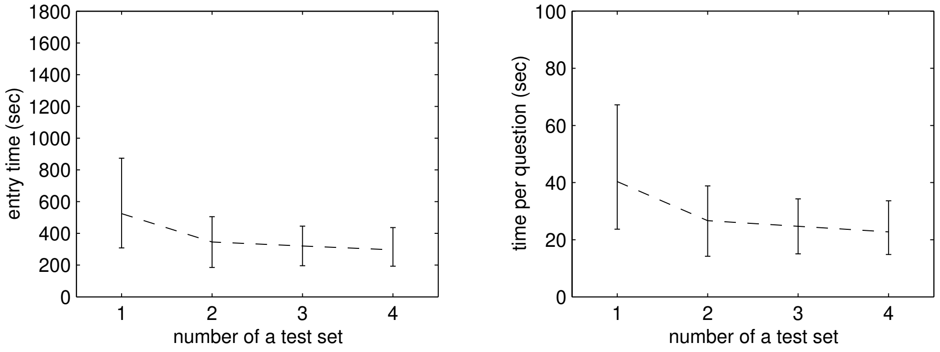

Fig. 11. Entry time for test sets (left) and the mean time per question for each set (right). We plot the average time (dashed lines) and the time of the fastest and slowest users (vertical bars).

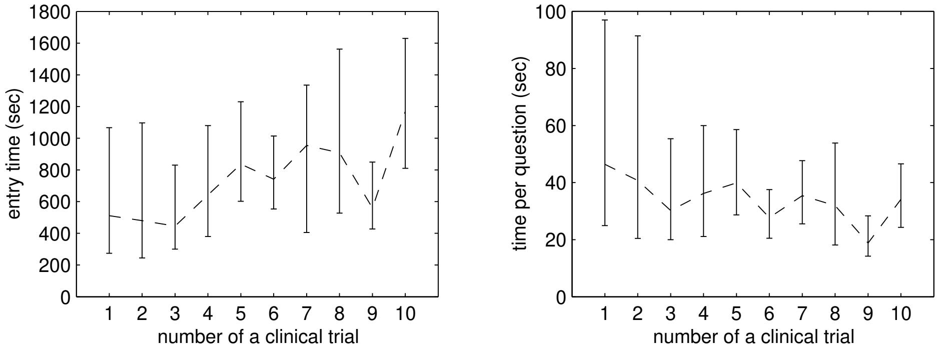

Fig. 12. Entry time for eligibility criteria. We show the average time for each clinical trial and the time per question (dashed lines), along with the performance of the fastest and slowest users (vertical bars).

selected criteria. If the eligibility conditions involve a more complex expression, the user has to click *"Combined question"* and then use the screen for composing logical expressions (Figure 10).

Entry time: We have run experiments with sixteen novice users, who had no prior experience with the interface. First, every user has entered four sets of medical tests; each set has included three tests and ten questions. Then, each user has added eligibility expressions for ten clinical trials used at the Moffitt Cancer Center; the number of questions in an eligibility expression has varied from ten to thirty-five.

We have measured the entry time for each test set and each eligibility expression. In Figure 11, we show the mean time for every test set and the time per question for the same sets. All users have entered the test sets in the same order, from 1 to 4; since they had no prior experience, their performance has improved during the experiment. In Figure 12, we give similar graphs for the entry of eligibility expressions.

The experiments have shown that novices can efficiently use the interface; they quickly learn its full functionality, and their learning curve flattens after about an hour. The average time per question is 31 seconds for the entry of medical tests and 37 seconds for eligibility criteria, which means that a user can enter all 150 cancer trials used at Moffitt in about two weeks.

V. Concluding Remarks

We have developed knowledge-acquisition tools for an agent that automatically assigns cancer patients to clinical trials. We have described the representation of eligibility criteria and a web-based interface for adding new trials. The experiments have shown that a user can enter a new trial in fifteen to thirty minutes. Novices can use the interface without prior instructions, and they reach their full speed after about an hour. Although cancer research at Moffitt has provided the motivation for this work, the agent is not limited to cancer, and we can use it for trials related to other diseases.

Acknowledgments: This work has been partially supported by the Breast Cancer Research Program of the U.S. Army Medical Research and Materiel Command under contract DAMD17-00-1-0244, and by the Moffitt Cancer Center.

References

[1] Sanjucta Bhanja, Lynn M. Fletcher, Lawrence O. Hall, Dmitry B. Goldgof, and Jeffrey P. Krischer. A qualitative expert system for clinical trial assignment. In *Proceedings of the Eleventh International Florida Artificial Intelligence Research Society Conference*, pages 84–88, 1998.

[2] Jim Blythe, Jihie Kim, Surya Ramachandran, and Yolanda Gil. An integrated environment for knowledge acquisition. In *Proceedings of the International Conference on Intelligent User Interfaces*, pages 13–20, 2001.

[3] Jacques Bouaud, Briggite Séroussi, Éric-Charles Antoine, Mary Gozy, David Khayat, and Jean-François Boisvieux. Hypertextual navigation operationalizing generic clinical practice guidelines for patient-specific therapeutic decisions. *Journal of the American Medical Informatics Association*, 5(suppl.):488–492, 1998.

[4] Jacques Bouaud, Briggite Séroussi, Éric-Charles Antoine, Laurent Zelek, and Marc Spielmann. Reusing ONCODOC, a guideline-based decision support system, across institutions: A successful experiment in sharing medical knowledge. In *Proceedings of the American Medical Informatics Association Annual Symposium*, volume 7, 2000.

[5] Bruce G. Buchanan and Edward H. Shortliffe. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Addison-Wesley, Reading, MS, 1984.

[6] Henrik Eriksson. Specification and generation of custom-tailored knowledge-acquisition tools. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, volume 1, pages 510–518, 1993.

[7] Yolanda Gil and Marcelo Tallis. A script-based approach to modifying knowledge bases. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, pages 377–383, 1997.

[8] Peter D. Johnson and Mark A. Musen. Development of a guideline authoring tool with PROTÉGÉ-II based on the DILEMMA generic protocol and guideline model. Technical Report SMI-96-0620, Stanford Medical Informatics Group, Stanford School of Medicine, 1996.

[9] Jihie Kim and Yolanda Gil. Acquiring problem-solving knowledge from end users: Putting interdependency models to the test. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence*, pages 223–229, 2000.

[10] Jihie Kim and Yolanda Gil. User studies of an interdependency-based interface for acquiring problem-solving knowledge. In *Proceedings of the International Conference on Intelligent User Interfaces*, pages 165–168, 2000.

[11] Princeton K. Kokku, Lawrence O. Hall, Dmitry B. Goldgof, Eugene Fink, and Jeffrey P. Krischer. A cost-effective agent for clinical trial assignment. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*, 2002.

[12] Cyrus Kotwall, Leo J. Mahoney, Robert E. Myers, and Linda Decoste. Reasons for non-entry in randomized clinical trials for breast cancer: A single institutional study. *Journal of Surgical Oncology*, 50:125–129, 1992.

[13] Sandra Marcus and John P. McDermott. SALT: A knowledge acquisition language for propose and revise systems. *Artificial Intelligence*, 39(1):1–37, 1989.

[14] Mark A. Musen. *Automated Generation of Model-Based Knowledge Acquisition Tools*. Morgan Kaufmann, San Mateo, CA, 1989.

[15] Mark A. Musen, David M. Combs, Edward H. Shortliffe, and Lawrence M. Fagan. OPAL: Toward the computer-aided design of oncology advice systems. In Perry L. Miller, editor, *Selected Topics in Medical Artificial Intelligence*, pages 166–180. Springer-Verlag, New York, NY, 1988.

[16] Mark A. Musen, Ray W. Fergerson, William E. Grosso, Natalya F. Noy, Monica Crubezy, and John H. Gennari. Component-based support for building knowledge-acquisition systems. In *Proceedings of the Conference on Intelligent Information Processing of the International Federation for Information Processing World Computer Congress*, 2000.

[17] Mark A. Musen, Samson W. Tu, Amar K. Das, and Yuval Shahar. EON: A component-based approach to automation of protocol-directed therapy. *Journal of the American Medical Informatics Association*, 3(6):367–388, 1996.

[18] Savvas Nikiforou. Selection of clinical trials: Knowledge representation and acquisition. Master's thesis, Department of Computer Science and Engineering, University of South Florida, 2002.

[19] Lucila Ohno-Machado, Eduardo Parra, Suzanne B. Henry, Samson W. Tu, and Mark A. Musen. AIDS^2: A decision-support tool for decreasing physicians' uncertainty regarding patient eligibility for HIV treatment protocols. In *Proceedings of the Seventeenth Annual Symposium on Computer Applications in Medical Care*, pages 429–433, 1993.

[20] Angel R. Puerta, John W. Egar, Samson W. Tu, and Mark A. Musen. A multiple-method knowledge-acquisition shell for the automatic generation of knowledge-acquisition tools. *Knowledge Acquisition*, 4(2):171–196, 1992.

[21] Thomas E. Rothenfluh, John H. Gennari, Henrik Eriksson, Angel R. Puerta, Samson W. Tu, and Mark A. Musen. Reusable ontologies, knowledge-acquisition tools, and performance systems: PROTÉGÉ-II solutions to Sisyphus-2. *International Journal of Human-Computer Studies*, 44(3–4):303–332, 1996.

[22] Briggite Séroussi, Jacques Bouaud, and Éric-Charles Antoine. Enhancing clinical practice guideline compliance by involving physicians in the decision process. In Werner Horn, Yuval Shahar, Greger Lindberg, Steen Andreassen, and Jeremy C. Wyatt, editors, *Artificial Intelligence in Medicine*, pages 76–85. Springer-Verlag, Berlin, Germany, 1999.

[23] Briggite Séroussi, Jacques Bouaud, and Éric-Charles Antoine. Users' evaluation of ONCODOC, a breast cancer therapeutic guideline delivered at the point of care. *Journal of the American Medical Informatics Association*, 6(5):384–389, 1999.

[24] Briggite Séroussi, Jacques Bouaud, and Éric-Charles Antoine. ONCODOC: A successful experiment of computer-supported guideline development and implementation in the treatment of breast cancer. *Artificial Intelligence in Medicine*, 22(1):43–64, 2001.

[25] Edward H. Shortliffe. *MYCIN: A Rule-Based Computer Program for Advising Physicians Regarding Antimicrobial Therapy Selection*. PhD thesis, Computer Science Department, Stanford University, 1974.

[26] Edward H. Shortliffe, Randall Davis, Stanton G. Axline, Bruce G. Buchanan, Cordell C. Green, and Stanley Cohen. Computer-based consultations in clinical therapeutics: Explanation and rule acquisition capabilities of the MYCIN system. *Computers and Biomedical Research*, 8:303–320, 1975.

[27] Marcelo Tallis. A script-based approach to modifying knowledge-based systems. In *Proceedings of the Tenth Innovative Applications of Artificial Intelligence Conference*, pages 1183–1195, 1998.

[28] Marcelo Tallis and Yolanda Gil. Designing scripts to guide users in modifying knowledge-based systems. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence*, pages 242–249, 1999.

[29] Marcelo Tallis, Jihie Kim, and Yolanda Gil. User studies of knowledge acquisition tools: Methodology and lessons learned. In *Proceedings of the Twelfth Workshop on Knowledge Acquisition, Modeling and Management*, 1999.

[30] Salim Yusuf, Peter Held, K. K. Teo, and Elizabeth R. Toretsky. Selection of patients for randomized controlled trials: Implications of wide or narrow eligibility criteria. *Statistics in Medicine*, 9:73–86, 1990.