IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re *Ex Parte* Reexamination of:

Patent No. 6,188,988

Control No.: Not Assigned

Inventors: D.W. Barry, C.S. Underwood, B.J. McCreedy, D.D. Hadden, J.L. Lucas

Issue Date: February 13, 2001
Application No. 09/523,532
Filing Date: March 10, 2000

For: SYSTEMS, METHODS AND COMPUTER PROGRAM PRODUCTS FOR GUIDING THE SELECTION OF THERAPEUTIC TREATMENT REGIMENS

Patent No. 6,081,786

Inventors: D.W. Barry, C.S. Underwood, B.J. McCreedy, D.D. Hadden, J.L. Lucas

Issue Date: June 27, 2000
Application No. 09/283,702
Filing Date: April 1, 1999

For: SYSTEMS, METHODS AND COMPUTER PROGRAM PRODUCTS FOR GUIDING THE SELECTION OF THERAPEUTIC TREATMENT REGIMENS

Mail Stop *Ex Parte* Reexam
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

DECLARATION OF EDWARD H. SHORTLIFFE, M.D., Ph.D.

I, Edward H. Shortliffe, declare as follows:

1. I have been asked to review (1) the patent specification and claims of U.S. Patents 6,081,786 ("'786") and 6,188,988 ("'988") (collectively, "the Barry patents"); and (2) certain pertinent prior art. The purpose of this declaration is to explain the contents of this prior art and its relevance to the Barry patent claims. I am making this statement voluntarily and without compensation because of my belief that the original patent applications, of which I had no knowledge until recently,[1] were based on claims that were not original and that there was substantial prior work — much of it by me — dating back to at least the early 1970s. Since the patents were filed in the late 1990s, there had been ample evidence of these observations in the scientific literature dating back at least 25 years. I say this with some authority since my own doctoral dissertation, first published in 1973 and completed in late 1975, was a therapeutic advisory system for infectious disease therapy that used expert knowledge encoded in rules and other data structures, arguably in multiple "knowledge bases" as the Barry patents use the term. Subsequent work by myself and others built upon that early research (known as the MYCIN System) and invalidate the claims for novelty and invention that were evidently the basis for the patent applications in question. I should add that, although we published the work extensively in the 1970s, 1980s, and 1990s, we viewed it as public-domain knowledge and methodology. We never sought to patent the work or the underlying methods and viewed it as contributing to a public good for the benefit of patients and health care. Our "reward" for our efforts was the acknowledgement from colleagues that the work was important and that the knowledge-based methods that we devised were worthy of adoption and expansion by others. I was accordingly particularly appalled when I learned that, not only were the Barry patents granted, but that

[1] Although I did much of my relevant work while a graduate student and then faculty member at Stanford University, I left to take a position elsewhere in 2000 and was not informed by the university when they subsequently entered into litigation and negotiations regarding the patents in question.

subsequently the patent owner has been arduously seeking to enforce that patent by demanding licensing royalties for work that, in my opinion, had long been in the public domain. This effort has had a stultifying effect on the use and refinement of the methods and hence has harmed those in the health care community and the patients that they serve. Disturbed by this situation, I have agreed to provide the analysis that follows in hopes that the Patent Office will agree that there was no merit to the original patent applications and that they should be disallowed in light of the information that I and others are providing to you.

I. Summary of My Opinions

2. In my opinion, the claims of the '786 and '988 patents represent no advance over the state of the art as it existed before 1997. I find no feature, functionality, or utility described in the '786 or '988 claims that was not previously developed and described both in my and my colleagues' work on the MYCIN and ONCOCIN expert systems in the 1970s and 80s, and numerous others' work on other expert systems designed to guide physicians in the selection of therapeutic treatment regimens. To the extent that the particular combination of features claimed in the '786 and '988 patents is arguably new, each claim was suggested by the literature. Moreover, successful practice of the claims would have been entirely predictable and routine given the state of the art at the time.

3. Specifically, I understand that the Patent Examiner has distinguished Barry's patent claims and the prior art on the basis that the Barry patents claim a system with "distinct" knowledge bases. This is a trivial distinction without a practical difference since multiple knowledge bases could be merged into a single entity and have long been separated into multiple representations largely for computational convenience and clarity. Furthermore, Barry does not explain why or how to implement an expert system with multiple distinct knowledge bases for suggesting therapeutic treatment regimens (there is much more technical detail in my own

writing on this precise subject). Indeed, the Barry patent application does not describe the specifics of any implementation of the claims. Critically, it does not even explain whether the "multiple knowledge bases" are intended to be physically distinct from each other (e.g., in different parts of the computer's memory or even on different machines) or merely conceptually distinct entities that may or may not be handled separately from a knowledge representation perspective.

4. Irrespective of whether the claims require physically or only conceptually distinct knowledge bases, either option would have been obvious and had been previously implemented in published systems. First, although in 1974 I described MYCIN as comprising "***a***" knowledge base referring only to the "productions" or rules for MYCIN, MYCIN's knowledge, which extended beyond simply the rule set, was separated and organized according to function and for computational convenience, and therefore was indeed divided into conceptually distinct units. Early publications on new work often leave out details of implementation that are not important to readers who are attempting to understand the novelty of a scientific effort, and in my earlier publications I referred to "a" knowledge base in MYCIN largely for didactic purposes. As I will describe below, subsequent writing definitely made it clear that there were other knowledge structures and representation methods used to encode knowledge in MYCIN and its successor systems, thereby using "multiple knowledge bases" in accordance with the distinction that appeared to be important in the opinion of the Patent Examiner. Second, the Barry patents do not describe any advantage or improvement over the prior art due to the separation of knowledge into three distinct knowledge bases. This is unsurprising because before 1997 it was a well-known design technique to utilize physically distinct "multiple knowledge bases," as the Patent Examiner uses the term, to implement expert decision-making systems, including systems for

therapeutic regimen selection. Organization of knowledge into discrete units was essentially a matter of convenience (both computational and conceptual), and was implemented in a predictable manner using well-known techniques.

II. Qualifications

5. I am both a physician and a computer scientist and am currently appointed Professor of Basic Medical Sciences and of Medicine, University of Arizona College of Medicine – Phoenix, in partnership with Arizona State University and Professor of Biomedical Informatics, Fulton School of Engineering, Arizona State University.

6. I am also President and CEO, American Medical Informatics Association.

7. I was the principal designer of the "MYCIN" computer-based expert system for selecting appropriate antibiotic therapeutic treatment regimens. MYCIN was the topic of my Ph.D. thesis ("*Shortliffe 1974*"), of a book I co-edited with Bruce Buchanan ("*Buchanan & Shortliffe 1984*"), as well as of numerous other research articles and reviews I have written over the years. I was also personally involved in the development of the ONCOCIN and T-HELPER expert systems, both of which predated the patents in question and also used multiple knowledge bases and varying knowledge representation methods.

8. I am the principal author and editor of the leading textbook in the field of biomedical informatics.[2] I am also Editor-in-Chief of the *Journal of Biomedical Informatics* (Elsevier)

9. I have published over 300 articles on biomedical informatics and expert systems both generally, and specifically in medicine. My CV is attached.

---

[2] Shortliffe and Cimino., Biomedical Informatics: Computer Applications in Health Care and Biomedicine. New York: Springer (3rd edition, 2006). Earlier editions: Shortliffe and Perreault, Medical Informatics: Computer Applications in Health Care, Reading, MA:Addison Wesley (1st edition, 1990) and New York: Springer (2nd edition, 2000).

III. Pending Reexamination Prosecution of the Barry Patents

10. I have been informed that both the '786 and '988 patents are the subject of ongoing patent reexamination.

11. I have been informed that the patent examiner for the reexamination of these patents concluded that the three "knowledge bases" listed in claims 1, 23, and 45 of both the '786 and '988 patents must be "distinct" from each other.[3] The patent examiner, however, did not elaborate what is meant by "distinct knowledge bases," a phrase not used in the Barry patent application or stressed in the pertinent literature; the notion of "distinct" knowledge bases would not be viewed as particularly useful among scientists in describing the way in which they formulate the knowledge structures for expert advisory tools. It is largely an implementation detail, but one that can be very helpful in assuring computational efficiency and simplifying knowledge maintenance tasks.

IV. The Barry Patents Do Not Describe the Structure of the Claimed Knowledge Bases

12. I have been asked to consider the Examiner's determination that Barry's claimed knowledge bases must be distinct from each other in light of Barry's patent specification.

13. I understand the '786 and '988 patent specifications both arise from the same patent application (U.S. Ser. No. 09/283,702), and are identical in all pertinent respects. I understand that the relevant date for assessing whether work by others was prior art is no earlier than April 3, 1998.

14. The Barry patents relate to an expert system for suggesting treatment regimens for HIV infected patients. The goals and functionality of the system discussed in Barry are equivalent to numerous other systems that were available before 1997, including MYCIN. In

[3] 95/001,088 Order Granting Reexamination at 9; 95/001,088 Action Closing Prosecution at 6-7; 90/010,313 Notice of Intent to Issue Ex Parte Reexamination Certificate at 5-6.

other words, both the ***inputs*** of the system ((1) expert knowledge concerning available treatments, rules for selecting and ranking treatments based on patient data, tables of drugs specifying dosing guidelines, side effects, and contraindications, and advisory information concerning recommended treatments, and (2) patient-specific information) and the ***outputs*** of the system (ranked treatment recommendations and advisory information) are the same as many systems implemented before Barry's patent application.

15. Barry's patent application is extremely superficial concerning any specific implementation of the system discussed. First, no computer code is provided. Second, only one cursory flow chart (Fig. 1) and one cursory schematic (Fig. 2) of the system are provided. Fig. 3 concerning the set up of the computer hardware is so general as to be non-informative. Figs. 4-12 are merely screen shots of the software user interface which tell the reader nothing about how to implement the claimed systems. The patent claims do not address the user interface. Third, the text concerning the system itself[4] is at the same high level of generality as Figures 1 and 2. In other words, a skilled person who wished to implement a system as claimed in either of the Barry patents would have to either code the software "from scratch" or use previously available systems, such as EMYCIN, the domain independent derivative of MYCIN.

16. There is nothing in the Barry patent specification that shows ***how*** the three listed knowledge bases are distinct from each other. For example, where the Barry specification describes "exemplary rules" it does not even indicate in which knowledge base they would reside. Rather, it completely confuses the issue by suggesting that these exemplary rules "may comprise ***one or more*** knowledge bases:[5]

---

[4] '786 Patent columns 7, 8, 10-12.

[5] '786 Patent col. 11:34-60.

Exemplary rules which may comprise one or more knowledge bases according to the present invention are listed below in Table 3.

TABLE 3

Therapy initiation/change: Rules that provide information on therapy change or initiation
Boundary condition rules: Limits for values, intervals for values to be updated
Comment Data Aging rules: These rules warn the user that the data in certain fields is getting old and that the most current values in the system will be used.
Rules that filter therapies due to drug interactions in ARV drug combinations
Rules that filter therapies due to medical conditions
Rules that filter therapies due to genotypic mutations in patient's plasma HIV
Rules that filter therapies due to phenotypic sensitivity/resistance
Antiretroviral therapy ranking rules
General dosage rules
Solid dosage rule
Dosage modifications due to ARV-ARV drug combination
Dosage modification due to ARV-NonARV interaction
Dosage modification due to medical condition
Comment determined
General commentary rules
Commentaries added due to medical conditions
Commentaries added due to drug interactions
Commentaries added due to drug combination
Delivery size rules

Thus the Barry patent provides no specifics for the implementation of the claims, such as how the "knowledge bases" would be structured, or how they would interact with each other, the inference engine, or the patient information.

17. So, from reading the Barry patent, one can not determine the specific structure that is claimed by the recitation of three knowledge bases, or what it would mean for them to be "distinct" from each other. Indeed, the Barry patent expressly states that the individual knowledge bases may be combined together into "a single database":

> Note also that, while the knowledge bases and patient data 21-25 are shown as separate blocks, ***the knowledge bases and patient data 21-25 can be combined together (e.g., the expert rules and the 20 advisory information can be combined in a single database***).[6]

This suggests to me that Barry's claimed knowledge bases ***do not*** need to be physically distinct from each other, and further calls into question whether the knowledge bases even need be

6 '786 Patent col. 8:16-21 (emphasis added).

conceptually distinct. The implication, as I indicated earlier, is that the decision whether to use single or multiple knowledge bases is a matter of computational and organizational convenience. It is not the source of the program's power and is accordingly not a notion that would seem worthy of a patent, even if there had been no discussion of the same notions in the prior art.

18. There is nothing in the Barry patent application that suggests that the provision of multiple knowledge bases distinguishes Barry's work from the prior art. Nor is there anything in the Barry patent application that shows any significance for implementation of such a system with "distinct" knowledge bases for treatment regimen information, expert rules, and advisory information, as compared to a single knowledge base comprising the same knowledge. In other words, the Barry patent specification does not describe any advantage or unexpected properties resulting from distinct knowledge bases.

V. Prior Art Expert Systems which Utilized Multiple Knowledge Bases

19. Even assuming that the Barry Patent application truly teaches and claims the use of multiple physically distinct knowledge bases, this would have been obvious in 1998. In my opinion there were many sophisticated solutions concerning knowledge base design available before 1997. In fact, it was a common design principle to organize expert knowledge into distinct knowledge structures (or "knowledge bases" if you like) to implement expert decision-making systems. Indeed, numerous prior art therapy-selection systems ***explicitly*** divided system knowledge into multiple "knowledge bases" along the same lines of demarcation as Barry.

A. *Shortliffe 1974* and *Buchanan 1984*

1. MYCIN

20. I began development of the MYCIN system in the early 1970s. My early work is documented in my Ph.D. thesis ("*Shortliffe 1974*"). Over the next decade, I collaborated with

many colleagues studying, improving, and extending the MYCIN system. Much of that work is captured in the 1984 book I co-edited with Bruce Buchanan ("*Buchanan 1984*").[7]

21. Contrary to the Examiner's analysis, knowledge in MYCIN comprised multiple knowledge structures. In addition to production rules for evaluating and selecting a therapeutic treatment regimen, MYCIN also had tables of facts and specialized procedures:

> **Knowledge Base**
>
> The form of knowledge is assumed primarily to be situation-action rules and fact triples (with CF's). ***Other knowledge structures, such as tables of facts and specialized procedures, are included as well***. Since the knowledge base is indexed and is small relative to the rest of the program, the size of the knowledge base should not be a limiting factor for most problems. MYCIN's knowledge base of 450 rules ***and about 1000 additional facts (in tables)*** is the largest with which we have had experience, although ONCOCIN is almost that large and is growing rapidly.[8]

22. A figure from *Buchanan 1984*[9] illustrates how knowledge was divided between rules and other static knowledge structures it refers to as "General Factual Knowledge of domain":

[7] Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project (Bruce Buchanan and Edward Shortliffe eds, Addison Wesley Publishing) (1984). Note that in addition to co-editing this book, I also authored or co-authored many of its Chapters, including those that I cite herein.

[8] *Buchanan 1984* at 697(emphasis added). See also *id.* at 99-100; 464.

[9] *Buchanan 1984* at 339.

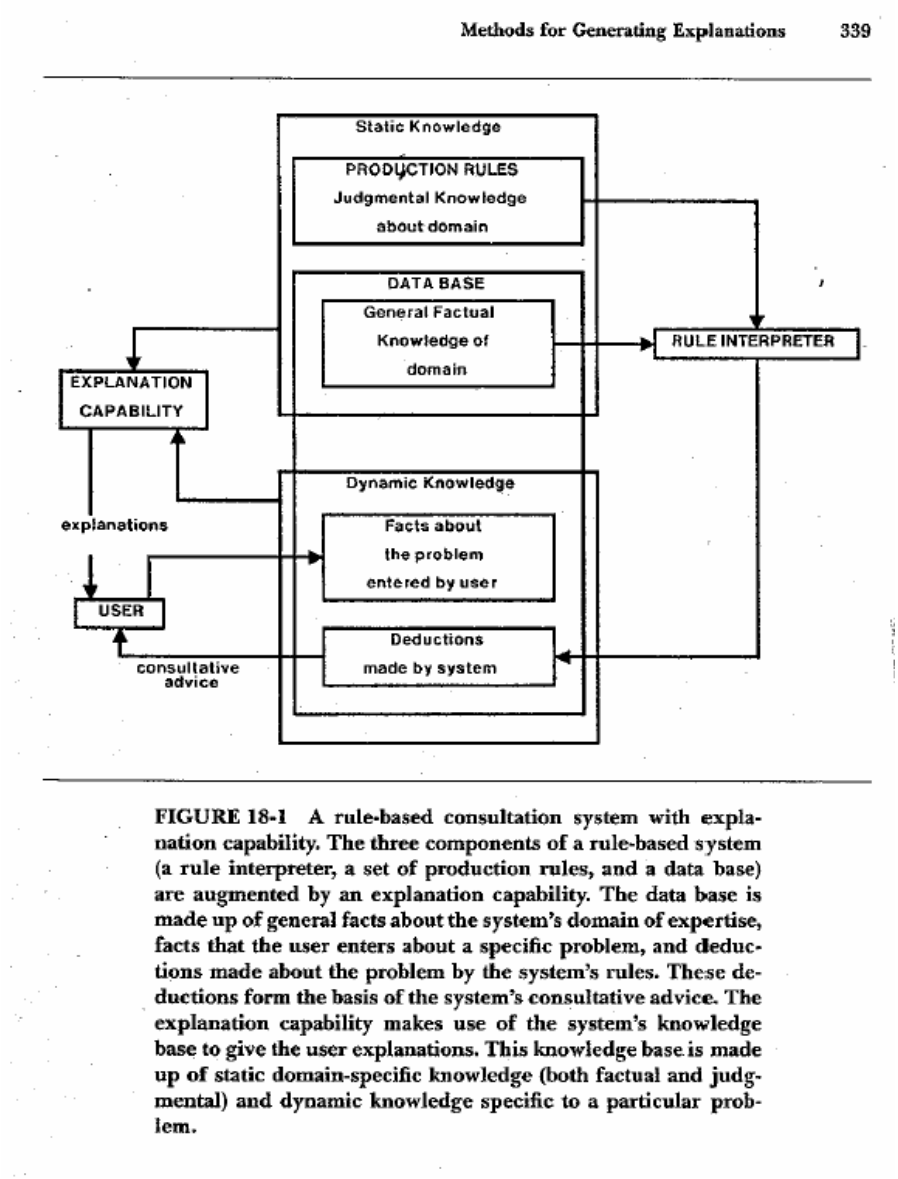

FIGURE 18-1 A rule-based consultation system with explanation capability. The three components of a rule-based system (a rule interpreter, a set of production rules, and a data base) are augmented by an explanation capability. The data base is made up of general facts about the system's domain of expertise, facts that the user enters about a specific problem, and deductions made about the problem by the system's rules. These deductions form the basis of the system's consultative advice. The explanation capability makes use of the system's knowledge base to give the user explanations. This knowledge base is made up of static domain-specific knowledge (both factual and judgmental) and dynamic knowledge specific to a particular problem.

Dating to the early 1980s, this figure is very similar to, and directly suggests creating distinctions between, different types of knowledge like the structure of the knowledge bases as set forth in Barry Figure 2 and the Barry patent claims from the late 1990s.

23. It would appear that MYCIN's general factual knowledge structures are analogous to the structures Barry defines as his first and third "knowledge bases."[10] MYCIN

[10] See my discussion of MYCIN's static knowledge structures, quoted below.

stored various data types in these structures, such as lists of infectious agents.[11] I see no principled differences between the additional knowledge structures I described in *Buchanan 1984* and either a list of therapeutic treatment regimens (the "first knowledge base") or a list of advisory information (the "third knowledge base").

24. I understand that the examiner has pointed to language I wrote in my thesis suggesting that MYCIN had a single knowledge base of expert rules to distinguish from Barry's multiple "distinct" knowledge bases. Given what little we know about how Barry's system was structured, I do not believe that there is a meaningful difference between the claimed system and my MYCIN system, especially given my subsequent publications that discuss separate knowledge structures in some detail. Rather I believe the difference to be largely a result of an evolving science without standards for all terminology — there simply was not a fixed definition of "knowledge base" or "bases" adopted by the field. Barry refers to knowledge structures such as lists of treatment regimens as a separate "knowledge base," and, although I encoded them separately in tables and other data structures that were not rule-based (see the figure reproduced earlier), I did not refer to them as separate knowledge bases but rather as components of a large and varied knowledge base. The distinction is meaningless. Barry does not identify any rules as part of his first or third knowledge bases. These structures are perfectly analogous to what we termed the "database" of general factual knowledge of domain." So while the terminology is different, the MYCIN system segregated knowledge into multiple distinct knowledge structures just as Barry claims.

25. Additionally, even if Barry claimed multiple knowledge bases comprising rules, MYCIN's production rules were also organized into distinct structures:

[11] *Buchanan 1984* at 99-100 ("'ORGANISMS' is the name of a linear list containing the names of all bacteria known to MYCIN.").

The 200 rules currently used by MYCIN are not explicitly linked in a decision tree or reasoning network. This feature is in keeping with our desire to keep system knowledge modular and manipulable. However, rules are subject to categorization in accordance with the context-types for which they are most appropriately invoked. For example, some rules deal with organisms, some with cultures, and still others deal solely with the patient. MYCIN's current rule categories are as follows (context-types to which they may be applied are enclosed in parentheses):

CULRULES	Rules that may be applied to any culture (CURCULS or PRIORCULS)
CURCULRULES	Rules that may be applied only to current cultures (CURCULS)
CURORGRULES	Rules that may be applied only to current organisms (CURORPS)
DRGRULES	Rules that may be applied to any antimicrobial agent that has been administered to combat a specific organism (CURDRUGS or PRIORDRGS)
OPRULES	Rules that may be applied to operative procedures (OPERS)
ORDERRULES	Rules that are used to order the list of possible therapeutic recommendations (POSSTHER)
ORGRULES	Rules that may be applied to any organism (CURORGS or PRIORORGS)
PATRULES	Rules that may be applied .to the patient (PERSON)
PDRGRULES	Rules that may be applied only to drugs given to combat prior organisms (PRIORDRGS)
PRCULRULES	Rules that may be applied only to prior cultures (PRIORCULS)
PRORGRULES	Rules that may be applied only to organism isolated from prior cultures (PRIORORGS)
THERULES	Rules that store information regarding drugs of choice (Section 5.4.1)[12]

[12] *Buchanan 1984* at 85.

Barry's first knowledge base of "therapeutic treatment regimens" appears to conceptually equivalent to "THERULES" — "Rules that store information regarding drugs of choice."[13]

26. I understand that the Examiner quoted some of the language set forth above to argue that my book "expressly warns against destroying the modularity of the knowledge base in order to keep the knowledge base manipulable" and thus in her view teaches against modification of MYCIN to the system Barry claims.[14] I disagree with the Examiner's interpretation. As can be seen by my full quote (reproduced above), I was only referring to the lack of "a decision tree or reasoning network" ***linking*** rules together, and therefore leaving the ***rules*** "modular and manipulable." In no way does this statement teach against ***dividing*** rules into separate structures, where they can still be "modular and manipulable," as is plain by the rest of the paragraph, which describes doing just that.

27. Because Barry does not delineate the structure of his various knowledge bases or their relation to each other, it may be that the distinctions among his system's knowledge that he intended in fact have the same characteristics as the MYCIN knowledge structures. Thus, at the level of abstraction provided by the Barry patent application, the reference to three knowledge bases is more of a difference in terminological choice than a practical difference when compared to a system with a "single" knowledge base (such as MYCIN) that is capable of storing and using multiple types of knowledge such as production rules and lists of data.

28. Lastly, in the Buchanan and Shortliffe book, we explicitly suggested that the next generation of expert systems for guiding the selection of therapeutic treatment regimens would be improved by further structuring and ***separating*** the systems' knowledge:

---

[13] *Buchanan 1984* at 85.

[14] 95/001,088 Action Closing Prosecution at 11.

> The success of MYCIN, which generally does not distinguish among these types of associations, demonstrates that it is possible to build a high-performance program within a sparse representation of homogenous rules (augmented with a few other knowledge structures). Nevertheless, limited experience with CENTAUR, WHEEZE, NEOMYCIN, and QN leads us to believe that ***the tasks of building, maintaining, and understanding the knowledge base will be easier if the types of knowledge are separated.*** This becomes especially pertinent during knowledge acquisition (as described in Part Three) and when teaching the knowledge base to students (Part Eight).[15]
>
> ***
>
> Modularity includes separation of individual "chunks" of knowledge from one another and from the program that interprets them. But it also implies a structuring of the knowledge that allows indexing from many perspectives. This facilitates editing, explanation, tutoring, and interpreting the individual chunks in ways that simple separation does not. In the case of MYCIN's rule-based structure, both the elements of data in a rule's premises and the elements of the rule's conclusion are separated and indexed. However, ***it is now clear that more structuring of a knowledge base than MYCIN supports will allow indexing chunks of knowledge still further***, for example to explain the strategies under which rules are interpreted or to explain the relationships among premise clauses.[16]

29. These passages from *Buchanan 1984* directly suggest the separation of knowledge into distinct knowledge bases that the Examiner has ascribed to Barry's claims. Our subsequent published work on ONCOCIN (see below) shows that we (among others) followed our own advice, well before the time of the Barry patent application.

2. ONCOCIN

30. The second system described in *Buchanan 1984* is ONCOCIN, a second-generation expert system for guiding therapeutic decisions in the context of oncology clinical

[15] *Buchanan & Shortliffe 1984* at 676 (emphasis added).

[16] *Buchanan 1984* at 670-671 (emphasis added).

trials. ONCOCIN used more organized structures to separate different forms of knowledge than did MYCIN.[17] Specifically, ONCOCIN employed four data structures in addition to rules:

> **35.4.2 Representation**
>
> ***
>
> ***Knowledge about the oncology domain is represented using five main data structures: contexts, parameters, data blocks, rules, and control blocks.*** In addition, we use a high-level description of each of these structures to serve as a template for guiding knowledge acquisition during the definition of individual instances.[18]

Each of the five data structures are described on pages 659-660 of *Buchanan.*

31. *Buchanan 1984* expressly shows that ONCOCIN stored its list of therapeutic treatment regimens in a different data structure than its knowledge bases of rules. Specifically, the ”control blocks” stored knowledge including protocols and chemotherapeutic procedures, i.e., the therapeutic treatment regimens of Barry’s “first knowledge base.” Control blocks were a separate data structure for storing procedural knowledge such as the steps in a protocol or chemotherapy regimen. [19] ONCOCIN is therefore another prior art implementation of an expert system that recommends treatment regimens where the system’s knowledge was separated into distinct knowledge bases, including at least the first and second knowledge bases claimed by Barry.

B. *Van Heijst 1994*

32. We were not the only workers to suggest that different types of knowledge in an expert system should be represented separately in systems for guiding selection of therapeutic treatment regimens. In a 1994 review article, van Heijst et al. explain that it was then “a generally accepted principle” that different types of knowledge be “represented separately.”

---

[17] See *Buchanan 1984* Chapter 35 “An Expert System for Oncology Protocol Management” by *Shortliffe et al.*

[18] *Buchanan 1984* at 659 (emphasis added).

[19] *Buchanan 1984* at 659-660.

> For example, the maintenance problems of first generation expert systems were mainly due to the fact that different types of knowledge, were mixed up in the knowledge base (Clancey, 1983). ***It is now a generally accepted principle that epistemologically different types of knowledge should be represented separately***.[20]

33. This passage from *van Heijst 1994* confirms my opinion that dividing expert system knowledge into "distinct knowledge bases" was well known in the art before 1998. And, since Barry's three knowledge bases are "epistemologically different," Barry's knowledge structure is directly suggested by this passage.

34. *Van Heijst 1994* also described blackboard systems, which are a very well-known architecture for knowledge based expert systems where knowledge is divided into independent modules:

> Blackboard systems consist of a blackboard and a number of knowledge sources (Nii, 1986). The blackboard is a central information repository that represents the state of the problem solving process. The knowledge sources are mechanisms that can inspect and, if their triggering conditions are satisfied, can modify the blackboard, thus driving problem solving. Knowledge sources are problem-solving modules that are specialized for particular types of problems whose internal representations and strategies are hidden for the other components of the architecture. Thus, knowledge sources may be considered as black boxes. The blackboard may be divided into different spaces, and knowledge sources may only be allowed to inspect or modify a limited number of spaces.[21]

Thus, blackboard systems are another example in the prior art of expert system implementations where knowledge is organized into "distinct knowledge bases."

C. *Chirico '444* Patent[22]

35. In a 1993 patent, Chirico described another system with multiple knowledge bases entitled "Expert System with a Plurality of Independent Knowledge Bases."

---

[20] Van Heijst et al., Foundations for a methodology for medical KBS development, Knowledge Acquisition (1994) 6, 395-434 at 396 (emphasis added).

[21] Van Heijst at 408.

[22] U.S. Patent 5,355,444, filed October 1993.

> The present invention overcomes the limitations and disadvantages of the prior art system by providing an expert system, that is, an inference engine with a plurality of stored rules based on the knowledge of an expert. The expert system of the present invention is variable in that ***it includes a plurality of rule systems (sometimes each rule system is referred to as a "knowledge base")*** and the information passed to the expert system includes a key factor (or a set of key factors), on which one of the rule systems (or knowledge base) is selected for use with each particular set of information.[23]

36. Unlike Barry, *Chirico '444* disclosed how and why multiple knowledge bases were implemented. *Chirico '444* clearly suggests organizing expert system knowledge into distinct knowledge bases like those claimed by Barry.

D. *Lau 1994*

37. The *Lau 1994* reference, entitled "A clinical decision support system prototype for cardiovascular intensive care" describes the design and implementation of an expert system for the therapeutic management of hypovolemic hypotension in the ICU setting.[24] Lau expressly states that its knowledge is divided into multiple knowledge bases:

> The knowledge elicited from the expert sources - target ranges, physiologic patterns, clinical conditions, intervention strategies, therapeutic drug-dosage responses, and therapeutic and reverse protocols - was formalized as the basis of the decision framework. The structure and usage of selected ***knowledge bases*** are briefly described below.[25]

38. *Lau 1994* describes no less than ***seven*** different knowledge bases in detail, including "Intervention choices,"[26] which directly corresponds to Barry's "first knowledge base," a "plurality of different therapeutic treatment regimens for said disease or medical condition."

---

[23] U.S. Patent 5,355,444 Col. 2:38-48.

[24] Francis Lau, A clinical decision support system prototype for cardiovascular intensive care." Int'l J. of Clinical Monitoring and Computing, 11:157-169, 1994.

[25] Lau at 160 (emphasis added).

[26] Lau at 160-163.

> *Intervention choices.* An intervention choice can be a specific diagnostic test or therapeutic agent. In this study, a set of therapeutic protocols was formalized to manage fluid replacement and bleeding. The protocols are invoked if 'reduce-volume-deficit' or 'stopactive- bleeding' was selected as the intervention strategy. The selection of the intervention choice, in this case an agent, from the protocols is contingent upon several factors. For bleeding, anti-coagulopathy agents such as protamine, cryoprecipitate, platelets and freshfrozen plasma are chosen according to PT, PTT and PLT results. In fluid replacement, the type of crystal loid or colloid used depends on the levels of serum Na, OSM, HCT and HB. The amount of fluid required is computed using the net-difference scoring and ranking algorithms. Critiquing rules were developed to determine the intervention choice when the patient is on vasoactive therapy or there is a change in the inotropic or vasoactive therapy.[27]

39. Moreover, at least the "Drug-dosage response tables" and "Net difference scoring" knowledge bases each directly correspond to Barry's "second knowledge base … of expert rules for evaluating and selecting a therapeutic treatment regimen."

E. *Suan 1990*

40. *Suan 1990* describes IMEX, a system for guiding malaria therapy that teaches a system with four modules (Initialization, Diagnosis, Treatment, and Drug Information):[28]

[27] Lau at 161-62.

[28] Suan 1990 at 362.

ONG LEAN SUAN

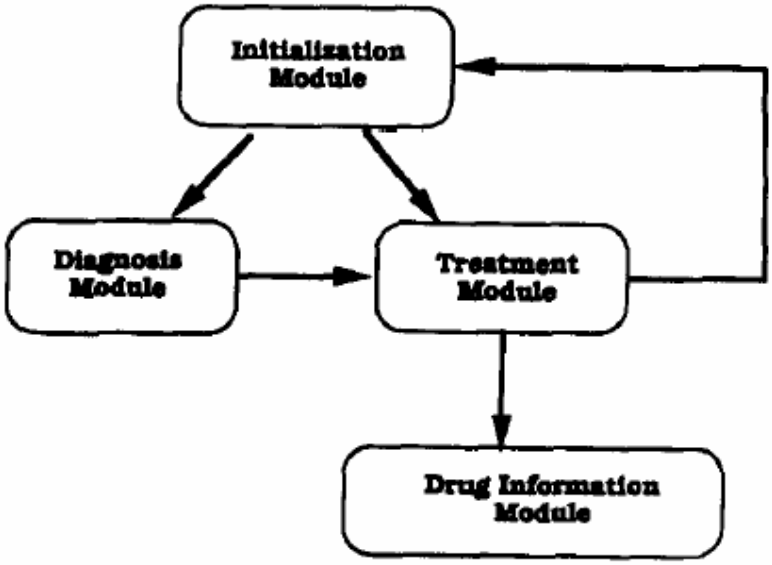

Fig. 1. Top-level modules of IMEX.

41. Each module within IMEX has its own knowledge base: "Each module has its own knowledge base which is organized as a hierarchy of rules, and passes intermediate conclusions to the other modules."[29]

42. The knowledge base for the Treatment Module directly corresponds to Barry's "second knowledge base … of expert rules for evaluating and selecting a therapeutic treatment regimen."

43. The knowledge base for the Drug Information Module directly corresponds to Barry's "third knowledge base comprising advisory information useful for the treatment of a patient with different constituents of said different therapeutic treatment regimens." As *Suan 1990* describes:

> At the end of the treatment module, the user can continue with another patient, end the session or ***ask for drug information***. ***The drug information module is a***

[29] Suan 1990 at 362.

database containing information on the common antimalarial drugs and allows the user to query for each type of drug, its history, chemistry, pharmacological effects, antimalarial actions, mechanism of antimalarial action, absorption, fate and excretion, preparation, routes of administration and dosage, toxicity/side effects and precautions/contraindications.[30]

F. *Windyga 1991*

44. *Windyga 1991*[31] describes the SETA system for guiding arrhythmia therapy in ICUs, whose system knowledge is divided into multiple "knowledge bases":

> From the expert's point of view, each type of arrhythmia and other relevant complications can be treated almost independently, suggesting the subsivision of the problem into ***several knowledge bases*** that can be triggered whenever necessary. Fig. 4 shows the general architecture of SETA. The decision board is the main component of the system. It triggers the required KB for the specific problem, according to ECG status, as detected arrhythmia and other information is entered through the dialogue block. Additionally, cardiac frequency and QRS width are requested by the decision board, as numerical information, if necessary.[32]

45. Windyga Figure 4 illustrates the multiple knowledge bases clearly:

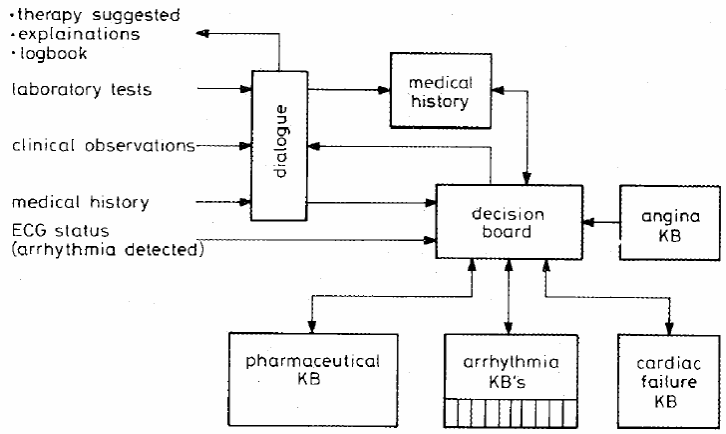

Fig. 4 *General architecture of SETA. Several knowledge bases (KBs) are connected to a decision board, which chooses the one that must be active at a particular moment. The dialogue block allows data input/output. The medical history block is a database about the patient being treated, keeping information that might be useful in future treatments. The ECG status is shown to be input directly to the decision board; in SETA this connection has been simulated by entering ECG measurements and interpretations through the dialogue block*

---

[30] Suan 1990 at 365-66 (emphasis added).

[31] P. Windyga, et al., "Knowledge-based approach to the management of serious arrhythmia in the CCU." Medical & Biological Engineering & Computing, 29:254-260 (1991).

[32] *Windyga 1991* at 257 (emphasis added).

46. Windyga's "pharmaceutical KB,"[33] directly corresponds to Barry's "third knowledge base" of "advisory information useful for the treatment of a patient with different constituents of said different therapeutic treatment regimens." As shown in *Windyga 1991* Figure 4, above, the system provides explanations, in addition to the "therapy suggested." Likewise, Windyga also describes a knowledge base for "alarm handling."[34]

47. The "angina KB," "arrhythmia KB's," and "cardiac failure KB" each directly correspond to Barry's "second knowledge base … of expert rules for evaluating and selecting a therapeutic treatment regimen."

48. Each of *Lau 1994*, *Suan 1990*, and *Windyga 1991* describe the function, structure, and interaction of the various knowledge bases described more thoroughly and clearly than does Barry.

49. In summary, long before the Barry patent application, workers in the field of biomedical informatics understood that the representation of knowledge in expert systems is inherently multi-faceted, simply because any expert system must utilize different types of knowledge in order to function. Moreover, as the field progressed, it became more and more sophisticated about organizing knowledge in order to make developing effective expert systems for guiding therapeutic treatment decisions. Many different means were developed to categorize, organize, and access knowledge in such expert systems, including systems that expressly divided knowledge into multiple distinct knowledge bases. I have provided examples of a number of prior art systems that are very similar to Barry's claims. Thus, by 1997, it would have been routine to utilize any combination of these various structures to achieve an expert system for suggesting therapeutic treatment regimens as claimed in the two Barry patents. The choice

[33] Windyga 1991 at 257, see especially. Fig. 4.

[34] Windyga at 259.

among the various among these various schemes for organizing system knowledge would be one of convenience, not invention. Thus, in my opinion, the Barry patent claims would have been obvious and the Patent Office should rescind them.

***

I declare that all of these statements are true and correct to the best of my knowledge and acknowledge that willful false statements are punishable by fine or imprisonment, or both pursuant to 18 U.S.C. § 1001.

Dated: July 7, 2009

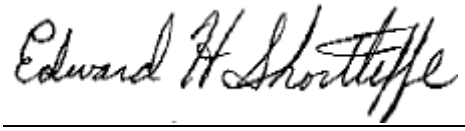

Edward H. Shortliffe, M.D., Ph.D.