A History of the INTERNIST-1 and Quick Medical Reference (QMR) Computer-Assisted Diagnosis Projects, with Lessons Learned

R. A. Miller
Donald A. B. Lindberg and Mary M. Lindberg University Professor of Biomedical Informatics
Department of Biomedical Informatics, Vanderbilt University Medical Center, Eskind Biomedical
Library, Nashville, Tennessee, USA

Summary

The INTERNIST-1/Quick Medical Reference (QMR) diagnostic decision support project spans four decades, from 1971-onward. This paper describes the history of the project and details insights gained of relevance to the general clinical and informatics communities.

Keywords

Medical imaging, imaging informatics, PACS, image analysis workstations

Introduction

The INTERNIST-1/Quick Medical Reference (QMR) computer-assisted diagnosis project spans four decades, from 1971 onward. This article provides a historical perspective on the project, emphasizing lessons learned (as "messages" below). Other published sources provide a broader perspective on the past half-century of work on computer-assisted diagnostic decision support (CDDS) systems, including the book by Morris F. Collen on the history of medical informatics per se [1], and Miller's 1994 *JAMIA* CDDS review with an extensive bibliography [2], updated in 2009 [3]. The references herein provide a complete list of publications by INTERNIST-1/QMR project participants [3-53], and a moderately complete list of publications about the project by external authors [54-69].

Message 1: Enlist clinical and informatics collaborators with expertise commensurate with the level required for achieving CDDS project objectives. Healthcare informatics projects must have personnel credible to both clinician-users and informaticians. Project members must interact with, and fully understand the ideas and concerns of, professional colleagues engaged in healthcare and information technology.

The founding long-term members of the INTERNIST-1 computer-assisted diagnosis project at the University of Pittsburgh included Jack D. Myers, MD, Harry E. Pople, Jr., PhD, and Randolph A. Miller, MD (initially as a medical student). Myers supplied world-class clinical expertise. Pople was from the onset a brilliant faculty computer scientist with established credentials in artificial intelligence. Miller, as the youngest participant, served as a "translator" who knew just enough medicine and computer science to explain concepts to Pople and Myers across disciplines; at times, discussion of his naïve questions helped to move the project forward. Miller substantially enhanced his education in both medicine and informatics by learning from senior project members.

Jack Duane Myers was born in New Brighton, PA (near Pittsburgh) in 1914. His family moved to Arizona when he was a child. After attending Stanford for his undergraduate and medical degrees, he undertook residency training in medicine at the Peter Bent Brigham (PBB) Hospital in Boston. There, under the brilliant Soma Weiss, MD, who was Chairman of Medicine, Myers became one of approximately 16 members of a 1939-1940 house staff contingent that produced fourteen future Department Chairs of Medicine.

Myers had an eidetic (photographic) memory, and by habit spent 20 hours per week in the library assimilating all that was known about pathophysiology and medicine in general. He often stated that, at that time, it was possible for an individual to master all established knowledge pertinent to internal medicine. Myers also learned through frequent interactions with his colleagues, who, as leaders in American internal medicine, were at the forefront of clinical research. He combined his dedica

tion to learning with keen observational skills and a highly logical intellect. Consequently, Myers earned a reputation as the "Sherlock Holmes" of American clinical diagnosis – the physician whom the best clinicians sought out when faced with challenging cases. Myers served as Chairman of Medicine at the University of Pittsburgh School of Medicine from 1955-1970. During the 1970's, he also had elected terms as Chairman of the National Board of Medical Examiners, President of the American College of Physicians, and Chairman of the American Board of Internal Medicine – some of the highest positions attainable in American internal medicine.

Harry E. Pople, Jr. was also a native of Pittsburgh. He attended M.I.T. as an undergraduate, where he developed an interest in the then-nascent fields of computer engineering and computer programming. He returned to Pittsburgh for his doctoral studies, entering the interdisciplinary Systems and Communications Sciences program in the Graduate School of Industrial Administration (GSIA) at Carnegie Mellon University (CMU). Edward Feigenbaum, later to become Chairman of Computer Science at Stanford and a contributor to Shortliffe's MYCIN [70] project, was also a graduate student in GSIA at the time. Pople obtained his PhD in 1968, working under psychology professor Allan Newell — who in the late 1950's was a founding father of Artificial Intelligence in the US, and in 1965 became the first Chairman of the Department of Computer Science at CMU. Newell, among many other accomplishments, pioneered application of rule-based production systems. Pople also worked closely on his dissertation project with Herbert Simon, another CMU psychology professor who shared with Newell a deep interest in understanding the mechanisms of human cognition. Simon had developed protocol analysis as a means to understand how chess masters and other experts solved difficult problems [71]. A decade later, Herbert Simon received

the Nobel Prize in Economics. Pople's dissertation involved developing a goal-oriented language for solving complex problems. After graduation from CMU with a GSIA PhD, Pople joined the faculty of the School of Business at the University of Pittsburgh.

During his fifteen-year term as Chair of Medicine in Pittsburgh, Myers became interested in the potential applicability of computers to clinical medicine. He was aware of the early work of G. Octo Barnett and colleagues on the COSTAR system [72] at Massachusetts General Hospital, and the early efforts at the National Board of Medical Examiners toward developing a computer-based licensure examination. After tiring of administrative responsibilities, Myers stepped down as Chair of Medicine in 1970, and decided to pursue, as a research focus, computer-assisted diagnosis. Dr. Gerhard Werner, then Chair of Pharmacology (and later Dean) at Pitt, introduced Myers to Pople. Werner and Pople had been collaborating on a project to use computer programs to model neural networks and their response to pain and analgesic medications. Myers was able to engage Pople in the quest to develop a clinical diagnostic system. Pople used the protocol analysis techniques developed by Simon to analyze Myer's logic as he "thought aloud" while solving challenging diagnostic cases. An initial approach, based on Pople's doctoral work, produced a system that reasoned from first principles. However, that approach was abandoned after a year of work, because searching the deep causal networks of pathophysiological knowledge gleaned from Myers was computationally inefficient. The program's execution times grew exponentially with the number of findings in a case, due to unconstrained searching. When there were independent case findings from more than one disease present, or a "red herring" finding not explained by any disease, the naïve search algorithm looking for a "unifying explanation" would not converge.

In March 1973, Randolph A. Miller joined Myers' diagnosis project during a research elective course in his second year of medical school. Miller was also a native of Pittsburgh and had majored in Physics at Princeton before returning to further his education at the University of Pittsburgh School of Medicine. Miller had learned to program computers in high school in the mid-1960s, and had continued programming in college. His initial role on the project was as a developer-coder in LISP and PDP-10 assembly language. Miller had joined the project just as Pople and Myers were abandoning their initial approach. Miller participated in critical discussions led by Pople and Myers that led, at Myers' suggestion, to a more empirical approach to diagnosis. The project was initially named DIALOG for diagnostic logic, but later renamed to INTERNIST-1 after discovery (through letters from lawyers) that Lockheed had previously created a system named DIALOG [4]. Miller significantly enhanced his medical education through participating in Myers-led seminars covering hundreds of disease topics related to knowledge base construction. In 1974-75, Miller took a one-year sabbatical from medical school to complete what was in essence a pre-doctoral fellowship in clinical informatics. During that time, he worked with Pople and Myers to develop the INTERNIST-1 knowledge base (KB) editor program, written in PDP-10 assembly language and LISP. Myers used the KB editor program for knowledge base construction and refinement. Myers would dictate at the computer terminal, typically while smoking a cigar, as his long-time assistant, Roseann Thomas, typed at the keyboard (with initially 110 baud then 300 baud connections to computers in Boston or Palo Alto). The KB editor program stored knowledge structures efficiently on disk as bit-compressed data structures. Before the advent of virtual memory operating systems, at Pople's request, Miller wrote an additional assembler-based virtual memory

123

A History of the INTERNIST-1 and Quick Medical Reference (QMR) Computer-Assisted Diagnosis Projects, with Lessons Learned

program, to feed "re-hydrated" data structures (from the compact KB editor-created disk representations) seamlessly as LISP S-expressions to Pople's LISP-based diagnostic algorithms.

In retrospect, a key enabling aspect of the project was Myers' seniority and national reputation. First, he had already completed a distinguished career in academic medicine, and was free to pursue any objectives he desired on the project, without concerns for a tenure clock or risk of loss of reputation if the project did not meet arbitrary objectives. Second, Myers' reputation drew the best and the brightest volunteers to the project, both from within the University of Pittsburgh School of Medicine and from other academic centers. Third, the degree of interest in what he was undertaking lended instant credibility and recognition to the project when it achieved early success. Similarly, Pople's background from CMU, and his pioneering work in artificial intelligence, attracted exceptionally capable students from computer science and the School of Business to the project.

Message 2: Develop a deep understanding of the problem addressed, employing appropriate expertise to develop a usable algorithm and related knowledge base, but do not go beyond what is pragmatically feasible. Conduct knowledge base construction and maintenance independently of the clinical decision support algorithms.

Clinical diagnosis is not a simple one-time matching of a clinical name (selected from a nosology, such as the International Classification of Diseases®) to a patient's case, typically comprised of a set of 'normal' and 'abnormal' finding descriptors representing the patient's history, physical examination, laboratory, imaging, and physiological function studies [31]. The process of medical diagnosis consists of "eliciting a portion of the patient's life story, an accounting of the patient's situation before the illness began, how the illness has manifested itself, how it has affected the life situation, and the pa-

tient's understanding of, and response to the illness" [31] as the disease process unfolds over time. "Diagnostic evaluation to determine the etiology of a patient's illness often involves sequentially eliciting, over potentially long drawn-out time spans, additional history, symptoms, physical exam signs, laboratory test results, and clinical image interpretations. For some illnesses, diagnosis may entail a 'therapeutic trial' to see if the patient responds to a specific intervention in a manner that is characteristic of a specific illness."[3] As a result, society and practitioners should consider computer-assisted diagnosis as only an adjunct to human clinicians' diagnosis acumen, and not a replacement for it [31].

That diagnostic dilemmas still comprise a significant challenge to practicing clinicians, and that clinical diagnosis remains a suitable domain for informatics research and development, was well documented in a 2008 AMIA conference that produced a 2009 supplemental issue of *Advances in Health Sciences Education* [73]. Correspondingly, a 2009 article by Newman-Toker and Pronovost labeled diagnostic errors as the "next frontier for patient safety" [74].

The goal of the INTERNIST-1 project was to address the problem of diagnosis within the broad context of general internal medicine: "Given a patient's initial history, results of a physical examination, and laboratory findings, INTERNIST-1 was designed to aid the physician with the patient's workup in order to make multiple and complex diagnoses."[6] The diagnostic algorithms and associated clinical diagnostic knowledge base (KB) could support expert consultations to construct and resolve differential diagnoses [6, 7, 9, 14, 34].

Several key early design aspects of the INTERNIST-1 knowledge base contributed to the project's longevity and the KB's lasting utility. The project's initial failed attempt to model the pathophysiological basis of disease at a "deep" level, using Pople's doctoral work as a diagnostic engine, influenced

the direction subsequently taken. Myers made a critical decision, contrary to expert systems philosophy prevalent at the time, whereby a "knowledge engineer" debriefed a "domain expert" to subjectively construct a system's KB. Despite having an eidetic memory, Myers decided to use the peer-reviewed published literature as the independent gold standard source of knowledge on which to build the INTERNIST-1 KB. Specifically, no finding would be listed in an INTERNIST-1 disease profile unless two separate groups had reported it in a reliable manner (i.e., the reporting authors had directly observed the patients and were not quoting other observers as the primary source).

While it may seem initially desirable to build a decision support system with "up to the minute" and "as detailed as possible" clinical information, doing so is perilous. The level of generally agreed upon knowledge breaks down when one attempts to represent the often-contrary "leading edge" advanced theories of the mechanisms of disease. The rate at which such theories evolve compounds the difficulty of maintaining such KBs for decision support purposes. This phenomenon was demonstrated by the Hepatitis Knowledge Base project in the 1970s [75]. Thus, it is important to select a project domain and knowledge representation scheme where the half-life of clinical knowledge is compatible with project resources. For example, the amount of effort required to construct and maintain an evidence-based repository of "best therapies" for all diseases and syndromes in internal medicine is substantially greater than the amount of effort required to characterize those findings reliably and verifiably reported to occur in patients with diseases seen by internists. The latter observation holds because the list of available drugs and the various therapies considered optimal for a given condition evolve rapidly, with a half-life on the order of a few years. By contrast, the list of potential findings that a given untreated illness causes in patients evolves slowly,

with a half-life measured in decades to centuries. The most rapidly changing aspects of disease knowledge at present typically involve "high technology": new immunological, genomic, and imaging test results. Based on several decades of INTERINST-1/QMR KB maintenance, the rate of change of the known findings related to a typical disease or syndrome is approximately 1-3 % per year.

A critical early project decision by Myers was to maintain the INTERINST-1 KB as an evidence-based, academic repository of diagnostic information, and to not "tune" the KB to enhance INTERNIST-1 system performance on specific test cases. On a related note, William Clancey, in developing the NEOMYCIN educational system from the original MYCIN [70] rule-based system, noted in retrospect the complexities that had been inadvertently introduced into MYCIN. The MYCIN rules had admixed procedural knowledge (what to do) with static domain knowledge (what was known as "established facts" in the domain of infectious disease) in a manner that complicated application of the rules to education [76]. The INTERNIST-1/ QMR project avoided this problem by basing the KB on the literature alone, in a manner completely decoupled from the diagnostic algorithms.

Miller, Pople and Myers authored a 1982 evaluation of INTERNIST-1's diagnostic performance on a series of difficult cases that had been published in the *New England Journal of Medicine* as "Clinicopathological Conference (CPC)" records [6]. In addition to demonstrating the impressive capabilities of the system, the evaluation identified several shortcomings of its approach [6, 9]. First, the knowledge base and diagnostic algorithms did not adequately represent disease and finding severity. Second, the temporal course of a patient's illness could not be fully described – the findings of a patient with acute appendicitis who first developed periumbilical pain, several hours later developed right lower quadrant pain, and who within hours devel-

oped a fever, would be identical in input form as the findings for a patient with Crohn's disease who had longstanding right lower quadrant pain and fever, who later developed periumbilical pain – with the exception that in the first case, the finding "abdomen pain acute" would be entered, and the second case, the finding "abdomen pain chronic" would be entered. Similarly, the 1982 evaluation indicated that the program could not "reason" anatomically regarding aspects of the patient's presentation – such as in the way that neurologists first consider anatomic sites of lesions consistent with a patient's findings, and then consider etiology.

A a consequence of the new insights into the shortcomings of INTERNIST-1, the project diverged into two separate components in the mid-1980s. Pople and colleagues, as capable computer scientists interested in the complexities of diagnostic reasoning, pursued new algorithms that could adequately address the temporal, spatial, and severity of illness issues not well handled by INTERNIST-1. Those algorithms were labeled "INTERNIST-II" and "CADUCEUS". However, soon thereafter, Myers became frustrated with the amount of effort required to build the expanded knowledge structures needed to support Pople's new approach. The clinical literature contains substantial variability in how it documents the findings that occur in an illness. There were essentially no common methods or templates for describing the time course or severities that a patient's illness might follow, and few published case series attempted to do so. Finally, one day, Myers observed that he had already dedicated a person-decade of effort to INTERNIST-1 knowledge base construction, and stated that he did not have another person-century of work left in his lifetime to build out Pople's new model for KB representation. As described below, Miller and Myers then took a different path, to develop Quick Medical Reference (QMR)® as a microcomputer-

based successor to INTERNIST-1 that preserved the knowledge base in its existing form.

Message 3: Quality CDDS KB construction requires meticulous, time-consuming effort. To ensure project longevity and value, KB construction should proceed in a scientifically reproducible manner. The project team must be able to construct, edit, and validate their system's clinical knowledge base within the scope of resources available to the project team. The project must develop, early on, a long-term plan for KB construction and maintenance.

Tables 1-4, reproduced with permission from [6] describe the basic building blocks of the INTERNIST-1/QMR KB consist of Disease Descriptions (Profiles), Diseases' Associated Disorders (Links), Findings' Differential Diagnoses, and "Properties" of Diseases and Findings. Based on Myers' conviction that disease profile content should be based on meticulous literature review, the INTERNIST-1/QMR KB construction efforts begun in 1973-74 comprise one of the earliest projects in "evidence-based medicine" [77]. To date, knowledge base construction and maintenance efforts have consumed approximately 50 person-years of work. More detailed descriptions of the INTERNIST-1/QMR KB, and the related INTERNIST-1 and QMR diagnostic algorithms, as well as its derivation, were published previously [3, 6, 7].

Myers' reputation attracted large numbers of clinicians-in-training to volunteer to develop disease profiles for the growing INTERNIST-1 KB. This activity served to introduce such students to the nascent field of clinical informatics [5, 26, 29]. Under Dr. Myers' supervision (and eventually Miller's supervision for the QMR KB in the 1990s), the procedural steps involved in the construction of an INTERNIST-1 / QMR KB Disease Profile by a "contributor" - a volunteer senior medical student on elective time, a resident physician on an elective rotation, or a faculty member at Pitt or elsewhere - included:

125

A History of the INTERNIST-1 and Quick Medical Reference (QMR) Computer-Assisted Diagnosis Projects, with Lessons Learned

1) Contributor selects a disease or clinical syndrome not yet described in the QMR knowledge base;
2) Contributor reviews the clinical literature on the diagnosis of the disorder, starting with general and subspecialty textbooks to obtain an overview of the topic, then selecting 50 to 250 relevant primary journal articles via bibliographic searching and following up the references of pertinent article. The review process for a disease typically required two to four weeks of fulltime effort. Descriptor findings for the disease came, whenever possible, from the existing INTERNIST-1/QMR KB findings (eventually numbering 5000) constructed by Myers previously; when necessary, new findings were created using syntax consistent with old finding names (i.e., mention site first, nature of finding next, and modifiers last – "abdomen pain right upper quadrant exacerbated by exercise").
3) Contributor compiles list of clinical abnormalities (findings) reliably and verifiably reported to occur in patients with the illness (i.e., reported by at least two separate groups independently). After 1988, finding entry was guided by the QMR-KAT program developed by Dario and Nunzia Giuse and Miller [46]; prior to that, contributors used handwritten spreadsheets.
4) Contributor consults clinical experts (local and/or national) to resolve inconsistencies or deficiencies uncovered during the literature review process.
5) Result of literature review is preliminary "disease profile" with an average of 85 findings per disease (range, 25 to 250 findings – see Table 4), plus an average of 8-10 "linked diagnoses" also identified (i.e., associated disorders that, for example, predispose to the current target disorder, are caused by the current target disorder, or co-occur with the target disorder). For example, an "acute myocardial infarc-

Table 1 Interpretation of Evoking Strengths

Evoking Strength	Interpretation
0	Nonspecific—manifestation occurs too commonly to be used to construct a differential diagnosis
1	Diagnosis is a rare or unusual cause of listed manifestation
2	Diagnosis causes a substantial minority of instances of listed manifestation
3	Diagnosis is the most common but not the overwhelming cause of listed manifestation
4	Diagnosis is the overwhelming cause of listed manifestation
5	Listed manifestation is pathognomic for the diagnosis

Table 2 Interpretation of Frequency Values

Frequency	Interpretation
1	Listed manifestation occurs rarely in the disease
2	Listed manifestation occurs in a substantial minority of cases of the disease
3	Listed manifestation occurs in roughly half the cases
4	Listed manifestation occurs in the substantial majority of cases
5	Listed manifestation occurs in essentially all cases—i.e., it is a prerequisite for the diagnosis

Table 3 Interpretation of Import Values

Import	Interpretation
1	Manifestation is usually unimportant, occurs commonly in normal persons, and is easily disregarded
2	Manifestation may be of importance, but can often be ignored; context is important
3	Manifestation is of medium importance, but may be an unreliable indicator of any specific disease
4	Manifestation is of high importance and can only rarely be disregarded as, for example, a false-positive result
5	Manifestation absolutely must be explained by one of the final diagnoses

tion" link would indicate that it predisposes to "cerebral embolism", a separately profiled disorder whose findings were not admixed with those of myocardial infarction.

For each finding entry in the "manifestations list" comprising a disease profile, two numbers were entered. The first, an "evoking strength" was analogous to a positive predictive value, given on a 0-5 scale (see Table 1 for interpretations). Myers and other senior clinicians were the primary sources of evoking strength information, although, after about half of the current knowledge base was in place, it became easier to "triangulate" assignment of evoking strengths by reviewing what values had been assigned to other diseases on the differential diagnosis list for a given finding. Table 2 explains the meaning of the second number, "frequency", which corresponded to the sensitivity (on a 1-5 scale) of the finding for the illness. The frequency numbers were firmly based on the literature reviews (i.e., how many cases across all literature reports of the disorder exhibited the finding).

Table 4 A Sample Manifestations List*
Reproduced with permission. From Miller RA, Pople HE Jr, Myers JD. INTERNIST-1, An Experimental Computer-based Diagnostic Consultant for General Internal Medicine. N Engl J Med 1982;307:468-76. Copyright © 1982, Massachusetts Medical Society. All rights reserved.

```
DISPLAY WHICH MANIFESTATION LIST?                          CHOLESTEROL BLOOD DECREASED                    2 2
ALCOHOLIC HEPATITIS                                        KETONURIA            1 2
                                                           PROTEINURIA          1 2
AGE 16 TO 25                  0 1                          SGOT 120TO 400       2 3
AGE 26 TO 55                  0 3                          SGOT 40TO 119        2 3
AGE GTR THAN 55       0 2                                  SGOT GTR THAN 400    1 2
ALCOHOL INGESTION RECENT HX   2 4                          UREA NITROGEN BLOOD LESS THAN 8                2 2
ALCOHOLISM CHRONIC HX              2 4                     UROBILINOGEN URINE ABSENT                      1 1
SEX FEMALE        0 2                                      UROBILINOGEN URINE INCREASED         2 4
SEX MALE          0 4                                      WBC 14000 TO 30000                             0 3
URINE DARK HX         1 3                                  WBC 4000 TO 13900 PERCENT NEUTROPHIL(S) INCREASED       0 3
WEIGHT LOSS GTR THAN 10 PERCENT     0 3                    WBC LESS THAN 4000                             1 1
ABDOMEN PAIN ACUTE           1 2                           ACTIVATED PARTIAL THROMBOPLASTIN TIME INCREASED   1 3
ABDOMEN PAIN COLICKY                  1 1                  ANTIBODY MITOCHONDRIAL                         1 1
ABDOMEN PAIN EPIGASTRIUM          1 2                      ANTIBODY SMOOTH MUSCLE                         2 3
ABDOMEN PAIN NON COLICKY          1 2                      BSP RETENTION INCREASED                        1 5
ABDOMEN PAIN RIGHT UPPER QUADRANT           1 3           ELECTROPHORESIS SERUM ALBUMIN DECREASED                  2 4
ANOREXIA                0 4                                ELECTROPHORESIS SERUM GAMMA GLOBULIN INCREASED  2 4
DIARRHEA ACUTE          1 2                                FACTOR VII PROCONVERTIN DECREASED       1 2
MYALGIA                 0 3                                LDH BLOOD INCREASED                 1 3
VOMITING RECENT         0 4                                MAGNESIUM BLOOD DECREASED    2 2
ABDOMEN BRUIT CONTINUOUS RIGHT UPPER QUADRANT     1 2      PROTHROMBIN      TIME INCREASED     2 3
ABDOMEN BRUIT SYSTOLIC RIGHT UPPER QUADRANT       1 2      SGPT 200 TO 600             1 2
ABDOMEN TENDERNESS RIGHT UPPER QUADRANT             2 4    SGPT 40TO 199              2 3
CONJUNCTIVA AND/OR MOUTH PALLOR                     1 2    SGPT GTR THAN 600    1 1
FECES LIGHT COLORED  1 2                                   LIVER BIOPSY BILE PLUGGING           1 2
FEVER                0 4                                   LIVER BIOPSY FATTY METAMORPHOSIS                      2 4
HAND(S) DUPUYTRENS CONTRACTURE(S)                  1 2     LIVER BIOPSY FOCAL NECROSIS AND INFLAMMATION       2 5
JAUNDICE             1 3                                   LIVER BIOPSY HEPATOCELLULAR NECROSIS MARKED       2 3
LEG(S) EDEMA BILATERAL SLIGHT OR MODERATE 1 2             LIVER BIOPSY MALLORY BODIES                           3 3
LIVER      ENLARGED MASSIVE              1 2              LIVER BIOPSY PERIPORTAL FIBROSIS MILD                 1 3
LIVER      ENLARGED MODERATE             1 3              LIVER BIOPSY PERIPORTAL INFILTRATION NEUTROPHIL(S)  3 5
LIVER ENLARGED SLIGHT            1 2                       LIVER BIOPSY PERIPORTAL INFILTRATION ROUND CELL(S)  1 2
PAROTID GLAND(S) ENLARGED            1 2                   LIVER BIOPSY SMALL BILE DUCT(S) PROMINENT          1 2
SKIN PALLOR GENERALIZED              0 2
SKIN     PALMAR ERYTHEMA             1 3                   LINKS FOR ALCOHOLIC HEPATITIS:
SKIN SPIDER ANGIOMATA                2 3                   Predisposes to MALLORY WEISS SYNDROME       1 1
SKIN TELANGIECTASIA                  1 1                   Causes SINUSOIDAL OR POSTSINUSOIDAL PORTAL HYPERTENSION   1 2
ALKLAINE PHOSPHATASE BLOOD GTR THAN 2 TIMES NORMAL1 2     Causes HEPATIC ENCEPHALOPATHY              2 2
ALKALINE PHOSPHATASE BLOOD INCREASED NOT OVER 2 TIMES NORMAL   1 4   Causes RENAL FAILURE SECONDARY TO LIVER DISEASE (HEPATORENAL SYNDROME) 2 2
BILIRUBIN BLOOD CONJUGATED INCREASED    2 4               Coincident with PANCREATITIS ACUTE .         2 2
BILIRUBIN URINE      PRESENT             2 4               Precedes MICRONODAL CIRRHOSIS (LAENNEC) 2 3

Table continues next column
```

* The first number beside each manifestation is the
  evoking strength; the second number is the frequency

127

A History of the INTERNIST-1 and Quick Medical Reference (QMR) Computer-Assisted Diagnosis Projects, with Lessons Learned

6) Clinical members of the INTERNIST-1/QMR project team reviewed the profile in a seminar led by Dr. Myers. Particular attention was paid to consistency in terminology for new findings, comparing the new disease profile to similar existing profiles to uncover any gaps or aspects not covered in the new profile, reconciliation of disparate quantitative values, completeness of the new profile, and aspects of the profile requiring further literature review and/or input from experts (often based on suggestions from Myers, based on his eidetic memory and thorough knowledge of medicine).

In the early 1990's, Dr. Nunzia Giuse and colleagues established that Myers' approach to CDDS knowledge base construction could be made scientifically reproducible through use of a tool (the QMR Knowledge Acquisition Tool, QMR-KAT) that ensured that contributors adhered to "good practices" and enforced consistency along a number of axes [22, 23, 32, 35, 44, 45, 46, 47]. Demonstration of scientific reproducibility of KB construction methods is a critical step rarely undertaken in expert system development [43]. In Giuse's study, clinical experts at seven geographically distant sites received instructions on how to construct disease profiles for the INTERNIST-1/QMR KB using the computerized QMR-KAT tool. Each participant was given a standard "starter set" of reprints of 109 articles from primary literature, 1906-1987, based on a bibliographic search conducted earlier by project members. The study demonstrated that the disease profiles that the subjects developed for "perinephric abscess" were remarkably (and statistically significantly) similar – indicating that with standard methods and adequate clinical expertise, it was reasonable to expect arbitrary individuals to develop the same disease descriptions for the KB [43].

Message 4: Feedback is critical to any healthcare informatics project

Too often in clinical practice and in healthcare informatics, an "expert committee" defines a terminology system, practice guidelines, knowledge base contents, or the design of a clinical informatics application (collectively, "information resources") that then remains static, or undergoes minimal maintenance, until the committee reconvenes months to years later (if at all). Knowledge in both healthcare delivery and clinical informatics evolves rapidly. Information resources must keep current, or risk advocating incorrect, outdated practices. In addition, inadequately updated systems may frustrate users who must endure correctable but still imperfect system traits about which they have previously submitted "change requests". Incorporation of critical and timely feedback, preferably from the clinician-users or expert consumers on the "front lines" of using the information resources in patient care, must form an integral part of any information resource-centered project. Projects must collect such information systematically (from both log files and from person-to-person interviews of users experiencing problems, since neither alone gives a complete picture), organize it in a logical manner, assign resolution responsibility to appropriately knowledgeable personnel with fixed, realistic timelines for completion, and act upon problems in a timely manner. Users tolerate imperfect situations when they know that someone will listen to their concerns and act upon them.

Under Myers' guidance, maintenance of the INTERNIST-1 / QMR KB became an essentially daily process, involving the following procedural steps:

1) Once a new profile has been created (manually or after 1988, using QMR-KAT [46]), it is entered into the INTERNIST-1/QMR KB via the corresponding editor program. The first check after new disease profile entry is to compare a new disease profile to existing similar disease profiles to ascertain completeness of concepts covered. For example, one might compare a new profile for "systemic candidiasis" to "disseminated tuberculosis" and discover that, while the tuberculosis profile includes genitourinary tract findings, the candidiasis profile does not. Any such "gaps" discovered send the person(s) responsible for knowledge base maintenance back to the library to determine if reliable and verifiable reports exist in the peer reviewed literature that fill in the gaps. Similarly, persons responsible for KB maintenance would determine if any new findings created in the construction of a new disease profile also should be entered into other diseases' profiles (again based on literature review).

2) From that point, the disease profile undergoes alterations as needed using the KB editor program, in response to literature-based, adequately documented new items of information.

3) Each time the system incorrectly diagnoses a case with an adequately documented diagnosis (e.g., confirming biopsy, autopsy, or pathognomonic test results), the clinician responsible for knowledge base maintenance (typically Myers) would do an "autopsy on the system" to determine if gaps in the knowledge base were responsible for the suboptimal performance. Only if literature-based knowledge existed that "filled in a gap" would the knowledge base be updated.

4) Myers tested each new profile immediately upon creation using an artificial, blatantly "classic case" of the disease. The test case was constructed by entering the findings of the disease prominently mentioned in standard textbook descriptions of the condition. Any missed diagnoses were addressed via step 3 above.

5) The profile was subsequently tested using more difficult published CPC (clinicopathological conference) cases and local challenging cases, as they arose in the literature or in clinical practice. Any missed diagnoses were addressed via step 3 above. All

test cases were given names and stored, so that they could undergo re-analysis at later dates. Over the first two decades of the INTERNIST-1/QMR project, the number of cases analyzed in this manner numbered approximately 2000.

6) Every two to three years, the entire series of "classical" cases was again analyzed, to make certain that intervening changes in the KB did not adversely affect performance on the cases. The availability of a "standard battery" of test cases can be extremely valuable in determining the strengths and weaknesses of a new system, as well as in the improvement of an existing system – analogous to calibrating analytical machines with standard reagents for quality control. Any "gaps" uncovered underwent step 3 procedures listed above.

7) As new published information from review articles, case reports, clinical conferences, other literature sources, and routine patient care "caught the eye" of clinicians on the project, the information would undergo review to determine if independent reports of the new "facts" supported their addition to the KB.

8) Additionally, de novo systematic reviews of the literature on a specific disease occurred every five to ten years, or sooner if major new information about the disease appeared in the literature. There have been two basic modes of disease profile construction. For exceedingly rare disorders (such as primary myocardial sarcoma) or newly discovered disorders (such as Legionnaires' disease in 1976-77), there rarely if ever exist published review articles wherein a single individual or group collates findings across a series of cases that they have directly seen. In such circumstances, only isolated case reports of single individuals appear in the literature, and the job of the INTERNIST-1/QMR disease profiler is to then correlate the individual case reports as published by primary observers of the

patients. The other mode of disease profile construction is more straightforward review of large case series of patients with a given disease that have been reported by careful observers in the literature, e.g., articles titled "A review of the last 500 cases of disease X seen by group Y at institution Z". After a "new" disease has been in the literature for several years, if it is not rare, groups then begin to publish larger case series that more reliably define the spectrum of the disease than do isolated case reports. This occurred, for example, for Legionnaires' disease in the 1979-1980 time frame. When such detailed larger reports begin to appear, it is appropriate to go back and review the original disease profile systematically, "from scratch", as Myers did for Legionnaires' disease in 1980.

Message 5: When feedback suggests that the project is "stuck", go back to the drawing board and assess what is valuable and working and what is not. Retain important core aspects of the project and re-invent the rest to move forward.

As illustrated in the Appendix (reproduced with permission from [6]), an INTERNIST-1 case analysis typically consumed 60-90 minutes of a clinician-user's time – for case finding abstraction into the INTERNIST-1/QMR terminology of 5000 findings, to enter the findings into the computer system, and to watch as the program slowly worked its way through case work-up and resolution. Furthermore, only individuals with access to the SUMEX-AIM facility at Stanford University, which required NIH/DRR grant support to obtain an account, could directly run cases using INTERNIST-1.

While it was a great intellectual challenge to construct a diagnostic engine and supporting knowledge base that could take a complete case description and conduct a work up leading to potentially multiple diagnoses, project members came to realize that pragmatically, few clinicians lost all their diagnostic abilities when faced with a difficult case.

Miller recognized that the "Greek Oracle" model of diagnostic assistance was no longer tenable [9, 30, 31] and that ubiquitous microcomputers in every clinician's office might provide "diagnostic decision support for the masses".

One of Miller's early responsibilities on the project was to print out a copy of the INTERNIST-1 KB on a monthly basis for Myers' reference. This involved higher-speed FTP transfer of files from Stanford to CMU, electronic transfer at lower speed of the information to a local machine at the University of Pittsburgh and generation of the several-hundred page long printouts locally. After the advent of the IBM PC® in 1980, Miller worked out a mechanism to more efficiently transfer the INTERNIST-1 KB files to his PC from Stanford, and began to develop PC-based programs to display the KB contents for Myers in a more flexible manner, via electronic sorting, than the fixed printed format.

When the IBM PC-AT® Intel ® 286 chip, and the availability of Borland's Turbo Pascal programming environment enabled development of more powerful programs, Miller expanded the earlier PC-based electronic textbook display programs into a more complete set of INTERNIST-1-like diagnostic algorithms. With this change came a major transformation of the project, from INTERNIST-1 to Quick Medical Reference (QMR)® [15, 17, 19, 21, 26, 27, 28, 30, 37, 38, 40, 41, 57].

The goal in developing QMR was to address the question, "Given the existence of the authoritative INTERNIST-1/QMR knowledge base, how could one assist human clinicians with the types of diagnostic dilemmas that they commonly encountered?" Thus, QMR was developed as a toolkit that enabled clinicians to gain multiple perspectives on a case, quickly and efficiently (in seconds to a few minutes), rather than going through the lengthy and cumbersome process of a full INTERNIST-1 diagnostic consultation.

The QMR toolkit had components providing support at three levels: elec-

129

A History of the INTERNIST-1 and Quick Medical Reference (QMR) Computer-Assisted Diagnosis Projects, with Lessons Learned

tronic textbook reference, intermediate low-level consultant, and INTERNIST-1 level consultant. At the first level, unlike INTERNIST-1, QMR would simply display, within seconds, portions of the INTERNIST-1/QMR KB, sorted appropriately to answer a question such as, "What are the most common findings of acute intermittent porphyria", "What history and physical exam findings have the highest predictive diagnostic value for confirming a suspicion of systemic lupus?", "What diseases predispose to acute myocardial infarction?", or "What diseases have findings closely resembling the findings of acute gout?" The intermediate level of QMR involved combinatorial reasoning among knowledge base components. For example, the QMR Relationships function addressed the question, "Given a small number of presenting findings stated to be of concern to the clinician, what disease or set of closely related diseases best fits with the findings?" [41] For example, given the input of "fever", "heart murmur systolic ejection left sternal border", "splenomegaly slight", and "hematocrit blood less than 35", QMR would quickly produce a list of 135 ranked possibilities, headed by the single diagnosis of subacute bacterial endocarditis, but also including acute lymphocytic leukemia, and rheumatoid arthritis complicated by both Felty's syndrome and anemia of chronic disease. Conversely, QMR contained a "patient case simulator" that would randomly generate quasi-realistic patient cases for trainees to diagnose [10]. Another QMR function allowed the user to critique a user-suggested diagnosis with respect to the findings of an inputted case – with the system producing ranked lists of findings present and explained by the diagnosis, findings absent in the case but expected in the diagnosis, and findings not mentioned in the case that constituted good "next steps" for work-up of the case toward confirming the inputted diagnosis. Yet another QMR function, "screen for complications of a disease" allowed the user to input a diagnosis name, and the system would scan through all disorders "linked" in the KB to the index diagnosis, and produced a combined list of findings sensitive and/or specific for the potentially complicating disorders. So inputting "diabetes mellitus" would produce findings useful in screening for the separately profiled disorder, diabetic retinopathy (such as hard exudates and signs of proliferative retinopathy), combined with findings of the also separately profiled disorder, diabetic nephropathy (such as proteinuria), and findings of coronary artery disease, and so on [15, 17, 19, 24-28, 33, 37-38, 40-42]

Key to facilitating the growth of QMR was an influx of extremely capable new members of the project team. Fred E. ("Chip") Masarie, Jr., MD, had completed his undergraduate and medical education at the University of Oregon, and two years of surgery residency training at the University of Pittsburgh Medical Center, before deciding that he wanted greater intellectual challenges than the surgical lifestyle seemed to offer. He volunteered to work on the QMR project shortly after its inception, during an elective rotation, and soon thereafter decided to devote the rest of his career to clinical informatics. Just as Miller had developed the INTERNIST-1 KB Editor program early in that project's history, Masarie very capably developed the first QMR KB Editor program, and in 1985-86 became an equal partner with Miller in ongoing and future QMR code writing. Approximately one year later, Nunzia B. Giuse, MD and Dario A. Giuse, PhD (Dr. Ing) joined the project. Nunzia Giuse was not only the first woman to graduate from college from her small home town in northern Italy near Brescia, but she also went on to obtain an MD from Universita' degli Studi di in Brescia in 1985 and an MLS degree from the University of Pittsburgh. She is now one of the most accomplished Biomedical Library Directors in the USA. She contributed to the QMR project initially as a volunteer building disease profiles. Based on her skill, insights, and accomplishments she soon joined the Section of Medical Informatics faculty in Pittsburgh, and rose to the rank of Assistant Vice Chancellor and full Professor at Vanderbilt. Her careful studies of the means by which contributors added knowledge to the INTERNIST-1/QMR KB (including the mechanisms of their errors) resulted in a number of publications [22-23, 29, 32, 34-35, 43-47]. It also led to the development by Dr. Dario Giuse, in conjunction with Nunzia Giuse and Randolph Miller, of the QMR Knowledge Acquisition Tool (QMR-KAT), which systematized and regularized knowledge base building for the project. Dario Giuse had been such a highly talented PhD/Dr.Ing student at Politecnico di Milano in Italy that he was recruited upon graduation to join the faculty at Carnegie Mellon University in 1985, bringing his new bride Nunzia with him. Dario collaborated on a number of clinical informatics projects at Pitt, including the QMR Project, and later, as an Associate Professor at Vanderbilt, developed the Star electronic medical record system that radically improved healthcare delivery there. The QMR-KAT program that Dario Giuse developed recorded not only finding names and corresponding evoking strengths and frequencies for each disease profile entry, but also all the bibliographic and other sources of the information supporting the entry, including how well (on a 1-5 scale of authority) each citation supported the "fact" and with what quantitative relevance (e.g., 15/31 cases described exhibited the finding); QMR-KAT could average the later quantitative information across all bibliographic entries for a specific finding listing in a disease profile to provide a suggested frequency number on a 1-5 scale [46].

Message 6: evaluate both in vitro and in vivo

A "first principle" for CDDS system evaluation [41, 50] is that clinicians, not systems or evaluators, are the individuals who discover, characterize, and attempt to solve clinical diagnostic problems. The reformulation of INTERNIST-1 project to be-

come the QMR project had as a central idea that the intelligent clinician-user was the most important component of the computer-assisted diagnostisis process. Dr. Charles Friedman reflected on this basic principle in a 2009 JAMIA article [78].

Unfortunately, in clinical informatics, it is commonplace to evaluate decision support systems using artificially constructed cases, or alternatively abstracted findings from actual cases presented to CDS systems and clinicians (as a comparison group) in artificial settings. For patient safety (i.e., not exposing live patients to potentially dangerous CDS systems before the systems have been adequately studied and validated for clinical use), it is appropriate to conduct formative evaluation studies using artificial cases, or preferably actual cases in artificial settings, as was described above for INTERNIST-1. Such evaluations might be termed 'in vitro' studies because they involve artificial cases, artificial settings, or both. Stead, writing for colleagues on the National Library of Medicine's BLRC study section at the time, noted that the stage of system development should determine the appropriate form of clinical informatics system evaluation [79].

As a CDS system grows in maturity, the impetus becomes greater for conducting an 'in vivo' study – whereby clinicians use the system for the purpose for which it was intended (with adequate safeguards as backup) on the "front lines" during actual clinical care. It is only then that one can determine if the system interface is too "clunky" to permit real-time use by busy clinicians, or whether the advice offered by the CDS system is of insufficient quality to merit clinicians' spending time using it – as well as establishing metrics for how well the system performs its intended advice-giving task.

Message 7: Demonstrate practical utility of system globally; ongoing development is eventually no longer viewed as research with respect to funding agencies.

As a CDS project matures, it finally reaches a stage where to continue, it must take one of several possible paths, at a time when research funding (by NIH or other agencies) is no longer feasible since the project has reached a stage of "ongoing development" rather than pioneering innovation. Until new technologies or application arenas open new research funding opportunities for the project, to survive, the project must either: (a) remain a local program supported by the "home institution" based on its utility to that institution; (b) become distributed as an "open source" product that generates low-level maintenance fees for the original developers; or (c) undergo commercialization and distribution through a licensing or similar relationship with a vendor.

In 1990, with the blessing of project members and the University of Pittsburgh Office of Intellectual Property, Dr. Fred E. Masarie, Jr. left University employment to form a company that marketed QMR. A year or two later, the Hearst Corporation subsidiary, First Databank, purchased the nascent company and with it the license to market QMR. The University of Pittsburgh retained the intellectual property rights to the QMR KB and diagnostic programs, but granted a temporary exclusive license and permanent non-exclusive license to the company to market QMR.

Prior to commercialization of QMR, Myers and the project team were invited to give demonstrations of the system at the annual meeting of the American College of Physicians (ACP). Due to the popularity of those sessions, the project team was given the opportunity to distribute free copies of the pre-commercial QMR program (i.e., at a time when commercialization had not even been envisioned) to physicians attending the ACP meeting who registered for a 3 half-days long tutorial session at the annual ACP meeting. There, the project team learned that, a priori, what clinicians imagined that a diagnostic decision support program like QMR could do was very different than what the actual system capabilities were. Clinicians envisioned functions not possible in the system, and the system had actual functions not conceived as possible by the clinicians. The essence of training in clinical informatics is to remove such mismatches between expectations and abilities of a software application. It became clear at the ACP sessions that highly capable clinicians could not, and would not, productively use the system after only a one-hour long introductory lecture explaining and demonstrating system capabilities. The physicians had to sit down and use the system themselves, with one-on-one instruction from a knowledgeable, experienced QMR clinician-user. Unless the ACP trainees used the system, typing in entries themselves, solving problems of relevance to them, they would not subsequently become successful "happy camper" long-term users. It is Miller's hypothesis that the First Databank sales of QMR eventually faltered because the training for clinician-users did not adequately adapt them to the system's capabilities to sustain long-term usage.

Conclusion

Ultimately, the success and viability of a knowledge-based clinical decision support project and its product depend upon:

1) Systematic, evidence-based, reproducible, sustainable knowledge base creation and maintenance activities.
2) Well-respected, academically knowledgeable, experienced, long-term project leaders and bright, committed younger participants who collectively have expertise in clinical, information technology, and evaluation domains.
3) Focus on the enhancing clinical end-users' abilities in a manner that they find helpful and useful. Formative and summative evaluation strategies that measure whether the CDS system extends what unaided clinicians can themselves accomplish.
4) Feedback loops to guide project evolution.

131

A History of the INTERNIST-1 and Quick Medical Reference (QMR) Computer-Assisted Diagnosis Projects, with Lessons Learned

5) An evolving strategy to support long-term project objectives and activities

Acknowledgments

References

1. Collen MF. A History of Medical Informatics in the United States: 1950 to 1990. Bethesda, MD: Amer Med Informatics Assoc; 1995. 489 pages. ISBN 0-9647743-0-5.

2. Miller RA. Medical diagnostic decision support systems: past, present, and future. J Am Med Inform Assoc.1994;1:827.

3. Miller RA. Computer-assisted diagnostic decision support: history, challenges, and possible paths forward. Adv Health Sci Educ Theory Pract 2009 Sep;14 Suppl 1:89-106. Epub 2009 Aug 12. PMID: 19672686.

4. Pople HE, Myers JD, Miller RA. DIALOG: A model of diagnostic logic for internal medicine. In: Proc Fourth Internat Joint Conf Artif Intel. Cambridge, Massachusetts: MIT Artificial Intelligence Laboratory Publications;1975. p. 848-55.

5. Miller RA, Schaffner KF. The logic of problem-solving in clinical diagnosis: a course for second-year medical students. J Med Ed 1982;57:63-5.

6. Miller RA, Pople HE Jr, Myers JD. INTERNIST-1, An Experimental Computer-based Diagnostic Consultant for General Internal Medicine. N Engl J Med 1982;307:468-76.

7. Pople HE Jr. Heuristic methods for imposing structure on ill-structured problems: the structuring of medical diagnostics. In Szolovits P (ed): Artificial Intelligence in Medicine. Boulder, CO: Westview Press (AAAS Symposium Series, no. 51). 1982;119-190.

8. Miller RA. Diagnostic and Prognostic Decision making Systems: A Survey of Recent Developments in the United States. In: Koehler CO, Tautu P, Wagner G, editors. Der Beitrag der Informationsverarbeitung zum Fortschritt der Medizin. Medizinische Informatik und Statistik series, #50. Proceedings of the 28th Jahrestagung, Deutsche Gesellschaft fur Medizinische Dokumentation, Informatik und Statistik. Heidelberg, Germany: Springer Verlag; 1983.

9. Miller RA. INTERNIST-1/CADUCEUS: Problems Facing Expert Consultant Programs. Methods Inf Med 1984;23:9-14.

10. Miller RA. A Computer-based Patient Case Simulator. Clin Res 1984;32:651A.

11. First MB, Soffer LJ, Miller RA. QUICK (QUick Index to Caduceus Knowledge): Using the INTERNIST-1/CADUCEUS Knowledge Base As An Electronic Textbook of Medicine. Comp Biomed Res.1985 Apr;18(2):137-65.

12. Miller RA, Schaffner KF, Meisel A. Ethical and Legal Issues Related to the Use of Computer Programs in Clinical Medicine. Ann Intern Med 1985;102:529-36.

13. Miller, RA. INTERNIST-1, An Experimental Program for Consultation in General Internal Medicine: A Critical Review of its Approach to the Use of Laboratory Information for Diagnosis. In: Proc 32nd Meeting Japan Soc Clin Path. Matsumoto, Japan. September, 1985.

14. Masarie FE Jr, Miller RA, Myers JD. INTERNIST-1 PROPERTIES: Representing Common Sense and Good Medical Practice in a Computerized Medical Knowledge Base. Comp Biomed Res 1985;18:458 79.

15. Miller RA, Masarie FE Jr, Myers JD. "Quick Medical Reference" for diagnostic assistance. MD Computing 1986;3:34-48.

16. Heckerman D, Miller RA. Towards a better understanding of the INTERNIST-1 knowledge base. In: Salamon R, Blum B, Jorgenson M, editors. Proc MEDINFO 86. Washington, D.C; 1986. p. 22-6.

17. Miller RA, McNeil MA, Challinor S, Masarie FE Jr, Myers JD. Status Report: The INTERNIST-1 / Quick Medical Reference Project. West J Med. 1986; 145:816-22.

18. Miller RA. Computer based Diagnostic Decision making. Medical Care 1987;25(12):S148-S152.

19. Bankowitz RA, Blumenfeld BH, Miller RA. User variability in Abstracting and entering printed case histories with Quick Medical Reference (QMR). Proc 11ᵗʰ Annu Symp Comput Appl Med Care. Washington, DC: IEEE Press; 1987. p. 68-73.

20. Parker RC, Miller RA. Using causal knowledge to create simulated patient cases: The CPCS project as an extension of INTERNIST 1. Proc 11ᵗʰ Annu Symp Comput Appl Med Care. Washington, DC: IEEE Press; 1987. p. 473-80.

21. Lincoln M, Turner C, Hesse B, Miller RA. A comparison of Clustered Knowledge Structures in Iliad and in Quick Medical Reference. Proc 12ᵗʰ Annu Symp Comput Appl Med Care. Washington, DC: IEEE Press; 1988. p. 131-5.

22. Giuse NB, Giuse DA, Miller RA. Computer assisted multi-center creation of medical knowledge bases. Proc 12ᵗʰ Annu Symp Comput Appl Med Care. Washington, DC: IEEE Press; 1988. p. 583-90.

23. Giuse NB, Bankowitz RA, Giuse DA, Parker RC, Miller RA. Medical Knowledge Base Acquisition: The Role of the Expert Review Process in Disease Profile Construction. Proc 13th Annu Symp Comput Appl Med Care; 1989. p. 105-9.

24. Wagner MM, Bankowitz RA, McNeil MA, Challinor SM, Janosky JE, Miller RA. The diagnostic importance of the history and physical examination as determined by the use of a medical decision support system. Proc 13th Annu Symp Comput Appl Med Care; 1989. p. 139-44.

25. Bankowitz RA, McNeil MA, Challinor SM, Parker RC, Kapoor WN, Miller RA. A computer- assisted medical diagnostic consultation service: implementation and prospective evaluation of a prototype. Ann Intern Med 1989;110:824-32.

26. Miller RA, Masarie FE Jr. Use of the Quick Medical Reference (QMR) program as a tool for medical education. Methods Inf Med 1989;28:340-5.

27. Parker RC, Miller RA. Creation of a knowledge base adequate for simulating patient cases: adding deep knowledge to the INTERNIST-1/QMR Knowledge Base. Methods Inf Med 1989;28:346-51.

28. Bankowitz RB, McNeil MA, Challinor SM, Miller RA. Effect of a computer-assisted general medicine diagnostic consultation service on housestaff diagnostic strategy. Methods Inf Med 1989;28:352-6.

29. Giuse NB, Giuse DA, Miller RA. Medical knowledge base construction as a means of introducing students to medical informatics. Proc Internat Symp Med Informat Educ. IMIA:Victoria, BC. May 1989. p. 228-32.

30. Miller RA, Masarie FE Jr. The demise of the "Greek Oracle" model for medical diagnostic systems. Methods Inf Med 1990;29:1-2.

31. Miller RA. Why the standard view is standard: people, not machines, understand patients' problems. J Med Philos 1990;15:581-91.

32. Giuse DA, Giuse NB, Miller RA. Towards computer assisted maintenance of medical knowledge bases. Artif Intel Med 1990;2:21-33.

33. Mabry ME, Miller RA. Distinguishing drug toxicity syndromes from medical diseases: a QMR computer-based approach. Proc 14th Annu Symp Comput Appl Med Care. Washington, DC:IEEE Computer Society Press; 1990. p. 65-71.

34. Miller RA, Giuse NB. The medicine in medical informatics: medical knowledge bases. Acad Med 1991;66(1):15-7.

35. Giuse DA, Giuse NB, Bankowitz RA, Miller RA. Heuristic determination of quantitative data for knowledge acquisition in medicine. Comput Biomed Res 1991;24(3):261-72.

36. Shiffman S, Wu AW, Poon AD, Lane CD, Middleton B, Miller RA, et al. Building a speech interface to a medical diagnostic system. IEEE Expert 1991;6:41-50.

37. Bankowitz RA, Miller JK, Janosky J. A prospective analysis of inter-rater agreement between a physician and a physician's assistant in selecting QMR vocabulary terms. Proc Annu Symp Comput Appl Med Care; 1991. p. 609-11.

38. Berman L, Miller RA. Problem area formation as an element of computer-aided diagnosis: a comparison of two strategies within Quick Medical Reference (QMR). Methods Inf Med 1991;30(2):90-5.

39. Masarie FE Jr, Miller RA, Bouhaddou O, Giuse NB, Warner HR. An interlingua for electronic interchange of medical information: using frames to map between clinical vocabularies. Comp Biomed Res 1991;24(4):379-400.

40. Miller RA, Jamnback L, Giuse NB, Masarie FE Jr. Extending the capabilities of diagnostic decision support programs through links to bibliographic searching: addition of "Canned MeSH Logic" to the Quick Medical Reference (QMR) program for use with Grateful Med. Proc Annu Symp Comput Appl Med Care;1991. p. 150-5.

41. Miller RA, Masarie FE. The Quick Medical Reference (QMR) Relationships Function: description and evaluation of a simple, efficient "multiple diagnoses" algorithm. Proc MEDINFO 92. Geneva, Switzerland:IMIA. 1992:512-8.

42. DeCillis A, Giuse DA, Thayer EH, Giuse NB, Miller RA. Using speech recognition in medical knowledge acquisition. Proc MEDINFO 92. Geneva, Switzerland:IMIA;1992. p. 1323-9.

43. Giuse NB, Giuse DA, Miller RA, Bankowitz RA, Janosky JE, Davidoff F, et al. Evaluating Consensus Among Physicians in Medical Knowledge Base Construction. Methods Inf Med 1993; 32:137-45.

44. Giuse DA, Giuse NB, Miller RA. Consistency enforcement in medical knowledge base construction. Artif Intel Med 1993(5):245-52.

45. Giuse NB, Giuse DA, Bankowitz RA, Miller RA. Comparing contents of a knowledge base to traditional information sources. Proc Annu Symp Comput Appl Med Care; 1993. p. 626-30.

46. Giuse DA, Giuse NB, Miller RA. Knowledge base editing with QMR-KAT. Proc Annu Symp Comput Appl Med Care; 1993. p. 935.

47. Giuse DA, Giuse NB, Miller RA. Evaluation of Long-Term Maintenance of a Large Medical Knowledge Base. J Am Med Inform Assoc 1995;2:297-306.

48. Aliferis CF, Cooper GF, Miller RA, Buchanan BG, Bankowitz RA, Giuse NB. A temporal analysis of QMR. J Am Med Inform Assoc 1996;3:79-91.

49. Rassinoux A-M, Miller RA, Baud RH, Scherrer J-R. Modeling principles for QMR medical findings. Proc Annu Symp Comput Appl Med Care 1996:264-8.

50. Miller RA. Evaluating evaluations of medical diagnostic systems. J Am Med Inform Assoc 1996. Nov-Dec;3:429-31.

51. Miller, RA. A heuristic approach to the multiple diagnoses problem. In: Keravnou E, Garbay C, Baud R, Wyatt J, editors. Artificial Intelligence in Medicine, Proceedings of AIME 97. Lecture Notes in Artificial Intelligence Series, #1211. Springer:Berlin; 1997. p. 187-98.

52. Rassinoux A-M, Miller RA, Baud RH, Scherrer J-R. Compositional and enumerative designs for medical language representation. Proc Annu Symp Comput Appl Med Care; 1997. p. 620-4.

53. Rassinoux A-M, Miller RA, Baud RH, Scherrer J-R. Modeling concepts in medicine for medical language understanding. Methods Inf Med 1998 Nov;37(4-5):361-72.

54. Shwe MA, Middleton B, Heckerman DE, Henrion M, Horvitz EJ, Lehmann HP, et al. Probabilistic diagnosis using a reformulation of the INTERNIST-1/QMR knowledge base. I. The probabilistic model and inference algorithms. Methods Inf Med 1991 Oct;30(4):241-55.

55. Middleton B, Shwe MA, Heckerman DE, Henrion M, Horvitz EJ, Lehmann HP, et al. Probabilistic diagnosis using a reformulation of the INTERNIST-1/QMR knowledge base. II. Evaluation of diagnostic performance. Methods Inf Med 1991 Oct;30(4):256-67.

56. Shwe M, Cooper G. An empirical analysis of likelihood-weighting simulation on a large, multiply connected medical belief network. Comput Biomed Res 1991 Oct;24(5):453-75.

57. Shiffman S, Lane CD, Johnson KB, Fagan LM. The integration of a continuous-speech-recognition system with the QMR diagnostic program. Proc Annu Symp Comput Appl Med Care 1992:767-71.

58. Berner ES, Webster GD, Shugerman AA, Jackson JR, Algina J, Baker AL, et al. Performance of Four Computer-Based Diagnostic Systems. N Engl J Med 1994 Jun 23;330(25):1792-6.

59. van Ginneken AM, Liem EB, Moorman PW. Integrating QMR with a computer-based patient record. Proc Annu Symp Comput Appl Med Care;1993. p. 98-102.

60. Gozum ME. Emulating cognitive diagnostic skills without clinical experience: a report of medical students using Quick Medical Reference and Iliad in the diagnosis of difficult clinical cases. Proc Annu Symp Comput Appl Med Care;1994. p. 991.

61. Welford CR. A comprehensive computerized patient record with automated linkage to QMR. Proc Annu Symp Comput Appl Med Care; 1994. p. 814-8.

62. Wolfram DA. An appraisal of INTERNIST-I. Artif Intell Med 1995 Apr;7(2):93-116.

63. Berner ES, Jackson JR, Algina J. Relationships among performance scores of four diagnostic decision support systems. J Am Med Inform Assoc 1996 May-Jun;3(3):208-15.

64. Arene I, Ahmed W, Fox M, Barr CE, Fisher K. Evaluation of quick medical reference (QMR) as a teaching tool. MD Comput 1998 Sep-Oct;15(5):323-6.

65. Lemaire JB, Schaefer JP, Martin LA, Faris P, Ainslie MD, Hull RD. Effectiveness of the Quick Medical Reference as a diagnostic tool. CMAJ 1999 Sep 21;161(6):725-8.

66. Berner ES, Maisiak RS, Cobbs CG, Taunton OD. Effects of a decision support system on physicians' diagnostic performance. J Am Med Inform Assoc 1999 Sep;6(5):420-7.

67. Berner ES, Maisiak RS. Influence of case and physician characteristics on perceptions of decision support systems. J Am Med Inform Assoc 1999 Sep-Oct;6(5):428-34.

68. Friedman CP, Elstein AS, Wolf FM, Murphy GC, Franz TM, Heckerling PS, et al. Enhancement of clinicians' diagnostic reasoning by computer-based consultation: a multisite study of 2 systems. JAMA 1999 Nov 17;282(19):1851-6.

69. Schaffner KF. Medical informatics and the concept of disease. Theor Med Bioeth 2000 Jan;21(1):85-101.

70. Shortliffe, E.H. Computer-based medical consultation: MYCIN. New York: Elsevier; 1976.

71. Newell A, Simon HA. Human Problem Solving. Prentice Hall: Englewood Cliffs, NJ. 1972.

72. Barnett GO. The application of computer-based medical record systems in ambulatory practice. N Engl J Med 1984;310(25):1643-50.

73. Berner ES, Graber ML, editors. Diagnostic Error in Medicine. Adv Health Sci Educ Theory Pract. 2009 Sep;14 Suppl 1:1-112. Epub 2009 Aug 12. ISSN 1382-4996 (Print) 1573-1677 (Online).

74. Newman-Toker DE, Pronovost PJ. Diagnostic Errors—The Next Frontier for Patient Safety, JAMA 2009;301(10):1060-2.

75. Bernstein LM, Siegel ER, Goldstein CM. The hepatitis knowledge base. A prototype information transfer system. Ann Intern Med 1980 Jul;93(1 Pt 2):169-81.

76. Clancey WJ, Leitsinger R. NEOMYCIN: reconfiguring a rule-based expert system for application to teaching. Proc 7th Internat Joint Conf Artific Intel. Vancouver, BC, Canada. San Francisco: Morgan Kaufmann Inc;1981. p. 829-36.

77. Sackett DL, Rosenberg WM, Gray JA, Haynes RB, Richardson WS. Evidence based medicine: what it is and what it isn't. BMJ.1996 Jan 13;312(7023):71-2.

78. Friedman CP. A "Fundamental Theorem" of Biomedical Informatics. J Am Med Inform Assoc 2009;16:169-70.

79. Stead WW, Haynes RB, Fuller S, Friedman CP, Travis LE, Beck JR, et al. Designing medical informatics research and library—resource projects to increase what is learned. J Am Med Inform Assoc 1994;1(1):28-33.

Correspondence to:
Randolph A. Miller, MD
Donald A.B. and Mary M. Lindberg University Professor.
Department of Biomedical Informatics,
Vanderbilt University Medical Center
Room B003C Eskind Biomedical Library
2209 Garland Avenue
Nashville, TN 37232-8340
E-mail: randolph.a.miller@vanderbilt.edu

133

A History of the INTERNIST-1 and Quick Medical Reference (QMR) Computer-Assisted Diagnosis Projects, with Lessons Learned

Appendix: A Sample Case Analysis

The transcript of an INTERNIST-I case analysis given below illustrates the oper- ation of the diagnostic programs. The case was taken from a CPC published in the *lournalin* 1969 [II]. The laboratory values are reported as measured in 1969. The bracketed paragraphs labeled "Comment" have been interpolated for clarification; they are not part of the actual consulta- tion. Places where the transcript has been abridged are indicated by an ellipsis.

```
INTERNIST-I      consultation
SUMEX-AIM        Version 15-May-81 07:31:39


ENTER CASE NAME: NEJM-CASE-30-1969-ADMISSION-l
   [Comment: Here the user enters the initial positive findings (present in the patient) and negative findings (absent). The
   Specialized INTERNIST-I vocabulary of some 3550 manifestations must be used in describing the case. The plus (+)
   prompt precedes each positive finding. entered by the user. Because INTERNIST-I has no mechanism for the repre-
   sentation of time, all findings have been collapsed into a single list, independently of their order of appearance in the
   patient.]


INITIAL  POSITIVE  MANIFESTATIONS:
+AGE GTR THAN 55
+ ARTHRITIS      HX
+ DEPRESSION
+ SEX FEMALE
+ THYROIDECTOMY HX
+ ULCER PEPTIC HX
+ URINE DARK HX
+ WEIGHT INCREASE RECENT HX
+ ANOREXIA
+ CHEST PAIN LATERAL EXACERBATION WITH BREATHING
+ CHEST PAIN LATERAL SHARP
+ DYSPNEA ABRUPT ONSET
+ ABDOMEN      DISTENTION
+ ABDOMEN FLUID WAVE
+ ASTERIXIS
+ FECES LIGHT COLORED
+ JAUNDICE
+ JOINT(S) PERIARTICULAR THICKENING
+ JOINT(S) RANGE OF MOTION DECREASED
+ LIVER ENLARGED MODERATE
+ PLEURAL FRICTION RUB
+ PULSE PRESSURE INCREASED
+ SKIN PALMAR ERYTHEMA
+ SKIN SPIDER ANGIOMATA
+ SPLENOMEGALY     SLIGHT
+ TACHYCARDIA
+ TACHYPNEA
+ THYROID ENLARGED ASYMMETRICAL
+ ALKALINE PHOSPHATASE BLOOD GTR THAN 2 TIMES NORMAL
+ BILIRUBIN BLOOD CONJUGATED INCREASED
+ BILIRUBIN URINE    PRESENT
+ FECES GUAIAC TEST POSITIVE
```

+ GLUCOSE BLOOD 130 TO 300
+ SGOT 120 TO 400
+ AMMONIA BLOOD INCREASED
+ ASCITIC FLUID PROTEIN 3 GRAM(S) PERCENT OR LESS
+ ASCITIC FLUID WBC 100 TO 500
+ ELECTROPHORESIS SERUM ALBUMIN DECREASED
+ ELECTROPHORESIS SERUM BETA GLOBULIN INCREASED
+ ELECTROPHORESIS SERUM GAMMA GLOBULIN INCREASED
+ GLUCOSE TOLERANCE DECREASED
+ IMMUNOELECTROPHORESIS SERUM IGA INCREASED
+ IMMUNOELECTROPHORESIS SERUM IGM INCREASED
+ IMMUNOELECTROPHORESIS SERUM IGG INCREASED

+ LDH BLOOD INCREASED
+ LE TEST POSITIVE
+ LIVER RADIOISOTOPE SCAN IRREGULAR UPTAKE
+ PROTHROMBIN TIME INCREASED
+ RHEUMATOID FACTOR   POSITIVE
+ STOMACH BARIUM MEAL HIATAL HERNIA
+*DENY*
  [Comment: The key word DENY indicates to INTERNIST-1 that the user now wants to enter negative findings. The minus (-) prompt precedes each negative finding entered.]
INITIAL NEGATIVE FINDINGS:
- ALCOHOLISM CHRONIC HX
- DIARRHEA CHRONIC
- FECES BLACK TARRY
- FEVER
- PRESSURE VENOUS INCREASED ON INSPECTION
- HEMATOCRIT BLOOD LESS THAN 35
- UREA NITROGEN BLOOD 30 TO 59
- URIC ACID BLOOD INCREASED
- ASCITIC FLUID AMYLASE INCREASED
- ASCITIC FLUID CYTOLOGY POSITIVE
- ASCITIC FLUID LDH GTR THAN 500
- ESOPHAGUS BARIUM MEAL VARICES
- STOMACH BARIUM MEAL ULCER CRATER
- T3 RESIN UPTAKE INCREASED
- T4 TOTAL BLOOD INCREASED
-GO
  [Comment: The user enters GO to indicate that all relevant positive and negative findings have been entered. The INTERNIST-1 consultant programs guide the user through the rest of the diagnostic workup by asking questions. For clarity, all responses typed by the user from this point on will be in italics.]
DISREGARDING: DEPRESSION HX, WEIGHT INCREASE RECENT HX, CHEST PAIN LATERAL EXACERBA-TION WITH BREATHING, CHEST PAIN LATERAL SHARP, DYSPNEA ABRUPT ONSET, ABDOMEN FLUID WAVE, ASTERIXIS, JOINT(S) PERIARTICULAR THICKENING, PLEURAL FRICTION RUB, THYROID EN-LARGED ASYMMETRICAL, FECES GUAIAC TEST POSITIVE, GLUCOSE BLOOD 130 TO 300, AMMONIA BLOOD INCREASED, ASCITIC FLUID PROTEIN 3 GRAM(S) PERCENT OR LESS, ASCITIC FLUID WBC 100 TO 500, GLUCOSE TOLERANCE DECREASED
CONSIDERING: AGE GTR THAN 55, SEX FEMALE" URINE DARK HX, ANOREXIA, FECES LIGHT COLORED, JAUNDICE, LIVER ENLARGED MODERATE, SKIN PALMAR ERYTHEMA, SKIN SPIDER ANGIOMATA, SPLENOMEGALY SLIGHT, ALKALINE PHOSPHATASE BLOOD GTR THAN 2 TIMES NORMAL, BILIRUBIN BLOOD CONJUGATED INCREASED, BILIRUBIN URINE PRESENT, SGOT 120 TO 400, ELECTROPHORESIS SERUM ALBUMIN DECREASED, ELECTROPHORESIS SERUM GAMMA GLOBULIN INCREASED,

135

A History of the INTERNIST-1 and Quick Medical Reference (QMR) Computer-Assisted Diagnosis Projects, with Lessons Learned

IMMUNOELECTROPHORESIS SERUM IGA INCREASED, IMMUNOELECTROPHORESIS SERUM IGG IN-CREASED, IMMUNOELECTROPHORESIS SERUM IGM INCREASED, LDH BLOOD INCREASED, LE TEST POSITIVE, LIVER RADIOISOTOPE SCAN IRREGULAR UPTAKE, PROTHROMBIN TIME INCREASED, RHEU-MATOID FACTOR POSITIVE

DISCRIMINATE: HEPATITIS CHRONIC ACTIVE, BILIARY CIRRHOSIS PRIMARY

[Comment: At this point, INTERNIST-1 has constructed a master differential- diagnosis list, ordered its members with a scoring algorithm, and then focused attention on the most promising problem area, which appropriately con-tains two liver diseases. The DISREGARDING list consists of all findings that are inconsistent with the topmost (first-listed) diagnosis; the CONSIDERING list includes all findings explained by the topmost diagnosis (i.e., chronic active hepatitis). The number of plausible contenders in the problem area determines the strategy for questioning. Questions are asked to discriminate between the two diagnoses.]

Please Enter Findings of LIPID(S) BLOOD

GO

CHOLESTEROL BLOOD DECREASED?

N/A

[Comment: INTERNIST-1 asks questions in two forms. A question of the style "Please Enter Findings of ... " asks about a generic class of findings, such as serum lipids or chest x-ray films. The user enters specific (positive or negative) findings and when finished types *"GO."* If there are no positive or negative find- ings to enter, typing *"GO"* will cause the programs to ask the original single item that prompted the generic question. For example, an elevated cholesterol level would favor primary biliary cirrhosis over chronic active hepatitis, whereas the converse situation would apply if the cholesterol were decreased. Thus, INTERN- IST-I asked a generic question about serum lipids. The second form of question- ing is direct. INTERNIST-I simply prints a manifestation name followed by a question mark. The user can then respond by typing *"YES"* (present), *"NO"* (absent), or *"N/A"* (not available-either unknown or not performed). The published CPC did not mention the serum cholesterol value, so the user entered N/A.]

Please Enter Findings of INSPECTION AND PALPATION SKIN

GO

SKIN XANTHOMATA?

NO

Please Enter Findings of BLOOD COUNT AND SMEAR

GO

RBC TARGET(S)

NO

[Comment: INTERNIST-1 uses information about the type of each manifestation-history, symptoms, signs, and labo-ratory tests, both routine and two gra- dations of more expensive or invasive laboratory procedures. The types of manifestations determine which questions are asked. Although the greatest dis- criminating power for resolving a differential diagnosis often resides in an inva- sive test, such as a liver biopsy, INTERNIST-I has been programmed to exhaust the useful history and physical-examination findings before escalating its questioning to more expensive or dangerous laboratory tests.]

Please Enter Findings of KIDNEY FUNCTION TEST(S)

GO

UREA NITROGEN BLOOD 60 TO 100?

NO

UREA NITROGEN BLOOD GTR THAN lOO?

NO

Please Enter Findings of LIVER FUNCTION TEST(S)

GO

UROBILINOGEN URINE ABSENT?

N/A

[Comment: INTERNIST -I did not seek to investigate renal function itself in the passage above. It was programmed to recognize that the urinary urobilinogen value may be unreliable in the presence of marked azotemia, so it first asked about the blood urea nitrogen.]

CONCLUDE: HEPATITIS CHRONIC ACTIVE

[Comment: It is likely that INTERNIST-1 had nearly made the diagnosis of chronic active hepatitis before it asked

any questions. The negative answers favored chronic active hepatitis, and therefore INTERNIST-1 concluded the diagnosis. The programs now recycle, using the findings not explained by the concluded diagnosis.]

DISREGARDING: DEPRESSION HX, CHEST PAIN LATERAL EXACERBATION WITH BREATHING, CHEST PAIN LATERAL SHARP, DYSPNEA ABRUPT ONSET, ASTERIXIS, JOINT(S) PERIARTICULAR THICKENING, PLEURAL FRICTION RUB, THYROID ENLARGED ASYMMETRICAL, FECES GUAIAC TEST POSITIVE, GLUCOSE BLOOD 130 TO 300, AMMO- NIA BLOOD INCREASED, GLUCOSE TOLERANCE DECREASED

CONSIDERING: WEIGHT INCREASE RECENT HX, ABDOMEN DISTEN- TION, ABDOMEN FLUID WAVE, ASCITIC FLUID OBTAINED BY PARACENTESIS, ASCITIC FLUID WBC 100 TO 500, ASCITIC FLUID PRO-TEIN 3 GRAM(S) PERCENT OR LESS ...

CONCLUDE: TRANSUDA TIVE ASCITES

DISREGARDING: ...

CONSIDERING: ASTERIXIS, PULSE PRESSURE INCREASED, TACHY-

CARDIA, AMMONIA BLOOD INCREASED PURSUING: HEPATIC ENCEPHALOPATHY

[Comment: The links in the INTERNIST-1 data base between chronic active hepatitis and hepatic encephalopathy have resulted in a bonus weight's being given to hepatic encephalopathy here; previously, links had promoted the consideration of transudative ascites, since it can also be caused by chronic active hepatitis.)

CSF FLUID OBTAINED?

N/A

[Comment: Here INTERNIST-I was about to ask about the glutamine level in the cerebrospinal fluid. Since no lumbar puncture was performed, the result is not available.]

CONCLUDE: HEPATIC ENCEPHALOPATHY

[Comment: In the above situation, there were no diagnostically helpful tests remaining for INTERNIST-1 to ask. INTERNIST-1 has been programmed to relax its criteria for concluding a diagnosis when all useful lines of questioning have been blocked. Since INTERNIST-1 had been close to making the diagnosis of hepatic encephalopathy, the program now concludes the diagnosis. The case analysis was intentionally stopped at this point, because all relevant major diagnoses had been covered. Without such intervention, INTERNIST-1 would try to explain any remaining important findings, such as the arthritis and pleurisy.]

Historic Landmark Paper Selected by Randolph Miller for the 2010 IMIA Yearbook of Medical Informatics*

Pioneers like Babbage, Turing, and von Neumann foretold how computing machines would be constructed and programmed before the modern digital age. Similarly, G. Anthony Gorry, PhD, while a young faculty member at MIT, described the essential reasoning components prevalent in almost all computer-assisted diagnostic systems created since 1970: (a) an inference mechanism to go from observed findings to hypothesized diagnoses; (b) a test selection function, that based on one or more models of utility (e.g., financial cost, risk to the patient, information theoretic value), would determine the "best next question to ask" to help resolve the current differential diagnosis; (c) a pattern-sorting function that determined whether two competing diagnoses were potentially complementary (i.e., both might be present in the patient) or alternatives (i.e., only one of the two competitors was likely to be present). Gorry's 1968 Mathematical Biosciences paper is thus a classical contribution to biomedical informatics and computer-assisted diagnosis.

Gorry, GA. Strategies For Computer-Aided Diagnosis. Mathematical Biosciences 1968;2:293-318.