



(12) **United States Patent**  
**Rao et al.**

(10) **Patent No.:** **US 7,617,078 B2**  
(45) **Date of Patent:** **Nov. 10, 2009**

(54) **PATIENT DATA MINING**  
(75) Inventors: **R. Bharat Rao**, Berwyn, PA (US);  
**Sathyakama Sandilya**, Cranbury, NJ  
(US); **Christopher Jude Amies**, Walnut  
Creek, CA (US); **Radu Stefan**  
**Niculescu**, Pittsburgh, PA (US); **Arun**  
**Kumar Goel**, Colonia, NJ (US);  
**Thomas R. Warrick**, Wayne, PA (US)

5,307,262 A 4/1994 Ertel  
5,359,509 A 10/1994 Little et al.  
5,365,425 A 11/1994 Torma et al.  
5,508,912 A 4/1996 Schneiderman  
5,544,044 A 8/1996 Leatherman  
5,557,514 A 9/1996 Seare et al.  
5,619,991 A 4/1997 Sloane  
5,652,842 A 7/1997 Siegrist, Jr. et al.

(73) Assignee: **Siemens Medical Solutions USA, Inc.**,  
Malvern, PA (US)

(Continued)

**FOREIGN PATENT DOCUMENTS**

(\* ) Notice: Subject to any disclaimer, the term of this  
patent is extended or adjusted under 35  
U.S.C. 154(b) by 482 days.

DE 198 20 276 11/1999

(Continued)

**OTHER PUBLICATIONS**

(21) Appl. No.: **10/287,055**

(22) Filed: **Nov. 4, 2002**

"Machine Learning and Data Mining", Mitchell, Communications of  
the ACM, Nov. 1999, ACM, USA, Onlinel, vol. 42, No. 11, Nov.  
1999, pp. 30-36, retrieved from internet: [http://portal.acm.org/  
ft\\_gateway.cfm?](http://portal.acm.org/ft_gateway.cfm?)

(65) **Prior Publication Data**

US 2003/0120458 A1 Jun. 26, 2003

(Continued)

**Related U.S. Application Data**

*Primary Examiner*—Paul L Rodriguez  
*Assistant Examiner*—Andre Pierre Louis

(60) Provisional application No. 60/335,542, filed on Nov.  
2, 2001.

(57) **ABSTRACT**

(51) **Int. Cl.**  
**G06F 7/60** (2006.01)  
**G01V 3/00** (2006.01)  
**G01W 1/00** (2006.01)

The present invention provides a data mining framework for  
mining high-quality structured clinical information. The data  
mining framework includes a data miner that mines medical  
information from a computerized patient record (CPR) based  
on domain-specific knowledge contained in a knowledge  
base. The data miner includes components for extracting  
information from the CPR, combining all available evidence  
in a principled fashion over time, and drawing inferences  
from this combination process. The mined medical informa-  
tion is stored in a structured CPR which can be a data ware-  
house.

(52) **U.S. Cl.** ..... **703/2; 705/2; 705/3; 702/181**

(58) **Field of Classification Search** ..... **703/2;**  
**703/11; 707/6, 104.1; 600/301; 702/181;**  
**705/2, 3; 706/2**

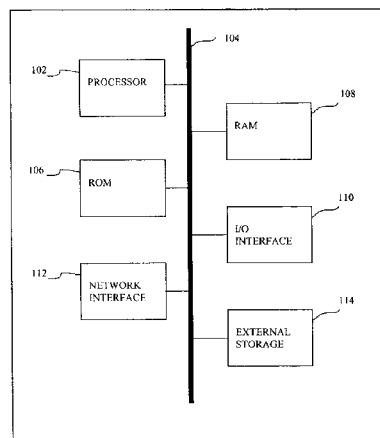
See application file for complete search history.

(56) **References Cited**

**U.S. PATENT DOCUMENTS**

4,946,679 A 8/1990 Thys-Jacobs

**65 Claims, 3 Drawing Sheets**



U.S. PATENT DOCUMENTS

5,657,255 A 8/1997 Fink et al.  
 5,664,109 A 9/1997 Johnson et al.  
 5,669,877 A 9/1997 Blomquist  
 5,706,441 A 1/1998 Lockwood  
 5,724,379 A 3/1998 Perkins et al.  
 5,724,573 A \* 3/1998 Agrawal et al. .... 707/6  
 5,737,539 A 4/1998 Edelson et al.  
 5,738,102 A 4/1998 Lemelson  
 5,811,437 A 9/1998 Singh et al.  
 5,832,450 A 11/1998 Myers et al.  
 5,835,897 A 11/1998 Dang  
 5,845,253 A 12/1998 Rensimer et al.  
 5,899,998 A 5/1999 McGauley et al.  
 5,903,889 A 5/1999 de la Huerga et al.  
 5,908,383 A 6/1999 Brynjestad  
 5,924,073 A 7/1999 Tyuluman et al.  
 5,924,074 A 7/1999 Evans  
 5,935,060 A 8/1999 Iliff  
 5,939,528 A 8/1999 Clardy et al.  
 6,039,688 A 3/2000 Douglas et al.  
 6,067,466 A 5/2000 Selker et al.  
 6,078,894 A 6/2000 Clawson et al.  
 6,081,786 A 6/2000 Barry et al.  
 6,083,693 A 7/2000 Nandabalan et al.  
 6,108,635 A 8/2000 Herren et al.  
 6,125,194 A 9/2000 Yeh et al.  
 6,128,620 A 10/2000 Pissanos et al.  
 6,139,494 A 10/2000 Cairnes  
 6,151,581 A 11/2000 Kraftson et al.  
 6,173,280 B1 \* 1/2001 Ramkumar et al. .... 707/6  
 6,196,970 B1 3/2001 Brown  
 6,212,526 B1 \* 4/2001 Chaudhuri et al. .... 707/102  
 6,253,186 B1 6/2001 Pendleton, Jr.  
 6,259,890 B1 7/2001 Driscoll et al.  
 6,266,645 B1 7/2001 Simpson  
 6,272,472 B1 8/2001 Danneels et al.  
 6,322,502 B1 11/2001 Schoenberg et al.  
 6,322,504 B1 11/2001 Kirshner  
 6,338,042 B1 1/2002 Paizis  
 6,381,576 B1 4/2002 Gilbert  
 6,468,210 B1 10/2002 Iliff  
 6,478,737 B2 \* 11/2002 Bardy ..... 600/301  
 6,484,144 B2 11/2002 Martin et al.  
 6,551,243 B2 4/2003 Bocionek et al.  
 6,551,266 B1 4/2003 Davis, Jr.  
 6,611,825 B1 \* 8/2003 Billheimer et al. .... 706/45  
 6,611,846 B1 8/2003 Stoodley  
 6,641,532 B2 11/2003 Iliff  
 6,645,959 B1 11/2003 Bakker-Arkema et al.  
 6,678,669 B2 1/2004 Lapointe et al.  
 6,754,655 B1 6/2004 Segal  
 6,802,810 B2 10/2004 Ciarniello et al.  
 6,804,656 B1 10/2004 Rosenfeld et al.  
 6,826,536 B1 11/2004 Forman  
 6,839,678 B1 1/2005 Schmidt et al.  
 6,903,194 B1 6/2005 Sato et al.  
 6,915,254 B1 7/2005 Heinze et al.  
 6,915,266 B1 7/2005 Saeed et al.  
 6,941,271 B1 9/2005 Soong  
 6,961,687 B1 \* 11/2005 Myers et al. .... 703/6  
 6,988,075 B1 1/2006 Hacker  
 7,058,658 B2 6/2006 Mentzer  
 7,130,457 B2 10/2006 Kaufman et al.  
 7,307,543 B2 12/2007 Rosenfeld et al.  
 2001/0011243 A1 8/2001 Dembo et al.  
 2001/0023419 A1 9/2001 Lapointe et al.  
 2001/0032195 A1 10/2001 Graichen et al.  
 2001/0041991 A1 11/2001 Segal et al.  
 2001/0051882 A1 12/2001 Murphy et al.  
 2002/0002474 A1 1/2002 Michelson et al.  
 2002/0010597 A1 1/2002 Mayer et al.

2002/0026322 A1 2/2002 Snowden et al.  
 2002/0026332 A1 2/2002 Snowden et al.  
 2002/0032581 A1 3/2002 Reitberg  
 2002/0035316 A1 3/2002 Drazen  
 2002/0077853 A1 6/2002 Boru et al.  
 2002/0082480 A1 6/2002 Riff et al.  
 2002/0087361 A1 7/2002 Benigno et al.  
 2002/0099570 A1 7/2002 Knight  
 2002/0123905 A1 9/2002 Goodroe et al.  
 2002/0138492 A1 \* 9/2002 Kil ..... 707/100  
 2002/0138524 A1 9/2002 Ingle et al.  
 2002/0143577 A1 10/2002 Shiffman et al.  
 2002/0165736 A1 11/2002 Tolle et al.  
 2002/0173990 A1 11/2002 Marasco  
 2002/0177759 A1 11/2002 Schoenberg et al.  
 2003/0028401 A1 2/2003 Kaufman et al.  
 2003/0046114 A1 3/2003 Davies et al.  
 2003/0050794 A1 3/2003 Keck  
 2003/0108938 A1 6/2003 Pickar et al.  
 2003/0120133 A1 6/2003 Rao et al.  
 2003/0120134 A1 6/2003 Rao et al.  
 2003/0120514 A1 6/2003 Rao et al.  
 2003/0125984 A1 7/2003 Rao et al.  
 2003/0125985 A1 7/2003 Rao et al.  
 2003/0125988 A1 7/2003 Rao et al.  
 2003/0126101 A1 7/2003 Rao et al.  
 2003/0130871 A1 7/2003 Rao et al.  
 2003/0135391 A1 7/2003 Edmundson et al.  
 2003/0208382 A1 11/2003 Westfall  
 2004/0078216 A1 4/2004 Toto  
 2004/0243586 A1 12/2004 Byers  
 2005/0187794 A1 8/2005 Kimak  
 2006/0064415 A1 \* 3/2006 Guyon et al. .... 707/6

FOREIGN PATENT DOCUMENTS

EP 0 596 247 A2 5/1994  
 EP 0 641 863 3/1995  
 EP 0 917 078 A1 10/1997  
 GB 2 332 544 A 6/1999  
 JP 11328073 A 11/1999  
 JP 2001297157 A 10/2001  
 WO WO 98/29790 7/1998  
 WO 98/39720 9/1998  
 WO WO 00/51054 8/2000  
 WO WO 00/69331 11/2000  
 WO WO 01/66007 9/2001  
 WO WO 01/78005 A2 10/2001  
 WO 01/82173 A1 11/2001

OTHER PUBLICATIONS

“Data Mining for the Enterprise”, Kleissner, System Sciences, 1998, Proceedings of the Thirty-First Hawaii International Conference on Kohala Coast, HI, USA Jan. 6-9, 1998, Los Alamitos, CA, USA, IEEE Comput. Soc. US, Jan. 6, 1998, pp. 295-304.  
 “Using Data Mining to Characterize DNA Mutations by Patient Clinical Features”, Evans et al., 1997 AMIA Annual Fall Symposium, Proceedings of 1997 AMIA Annual Fall Symposium the Emergence of Internetable Health Care Systems That Really Work, Nashville, TN, pp. 253-257.  
 “The Colorectal Cancer Recurrence Support (CARES) System”, Ong et al., Artificial Intelligence in Medicine, Nov. 1997, Elsevier, Netherlands, vol. 11, No. 3, pp. 175-188.  
 “Database System Support for Multidimensional Data Analysis in Environmental Epidemiology”, Kamp et al., Database Engineering and Applications Symposium, 1997, Ideas 97, Proceedings, International Montreal, Que, Canada, Aug. 25-27, 1997, Los Alamitos, CA, USA, IEEE Comput. Soc., US, pp. 180-188.  
 Grimes, “Structure, Models and Meaning, Is “Unstructured” Data Merely Unmodeled?”, Mar. 1, 2005, Intelligent Enterprise, <http://www.intelligententerprise.com/show/Article.jhtml?articleID=59301538>.

Berkus, "Unstructured Data" as an Oxymoron, Sep. 1, 2005, ITtoolbox Blogs, <http://blogs.ittoolbox.com/database/soup/archives/unstructured-data-as-an-oxymoron-5588>.

Larsen, "Fast and effective text mining using linear-time document clustering", 1999, ACM Press, Conference on Knowledge Discovery in Data, Proceedings of the fifth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 16-22.

Rao, "From Unstructured Data to Actionable Intelligence", IT Pro, Nov./Dec. 2003, pp. 29-35.

Mills. "Computer Technology of the Not-too-distant Future", Sep. 1993, Medical Laboratory Observer, vol. 25, No. 9, p. 78.

Duda "Pattern Classification", 2001. Jowl Wiley & Sons, Inc/. p. vii-xx, Chapter 1.

Hudson. "The feasibility using automated data to assess guideline-concordant care for schizophrenia", Dec. 4, 1999, Journal of Medical Systems, vol. 23, No. 4, pp. 299-307.

PR Newsline, Diabetes Health Management Award Honors Mayor Clinic's Aimmerman, Sep. 25, 2000.

Hudson, "CAATS and compliance", Apr. 1998, The Internal Auditor, vol. 55, No. 2, p. 25.

Waltz, "Information Understanding Integrating Data Fusion and Data Mining Processes", Circuits and Systems, 1998, ISCAS '98, Proceedings of the 1998 IEEE International Symposium on Monterey, CA USA, May 31-Jun 3, 1998, New York, NY, USA, IEEE, May 31 1998, pp. 553-556.

King et al., MEDUS/A: Distributing Database Management for Medical Research, Proceedings of Computer Networks Compcon 82, Sep. 20-23, 1982, pp. 553-556.

Boxwala et al., "Architecture for a Multipurpose Guideline Execution Engine", Proc. AMIA Symp 1999, pp. 701-705.

Guidance for Institutional Review Boards and Clinical Investigators 1998 Update, Sep. 1998, U.S. Food and Drug Administration.

Kassirer, "The Use and Abuse of Practice Profiles", Mar. 3, 1994, The New England Journal of Medicine, vol. 330:634-636.

Chen, "Do "America's Best hospitals" Perform Better for Acute Myocardial Infarctions?", Jan. 28, 1999, The New England Journal of Medicine, vol. 340, No. 4:286-292.

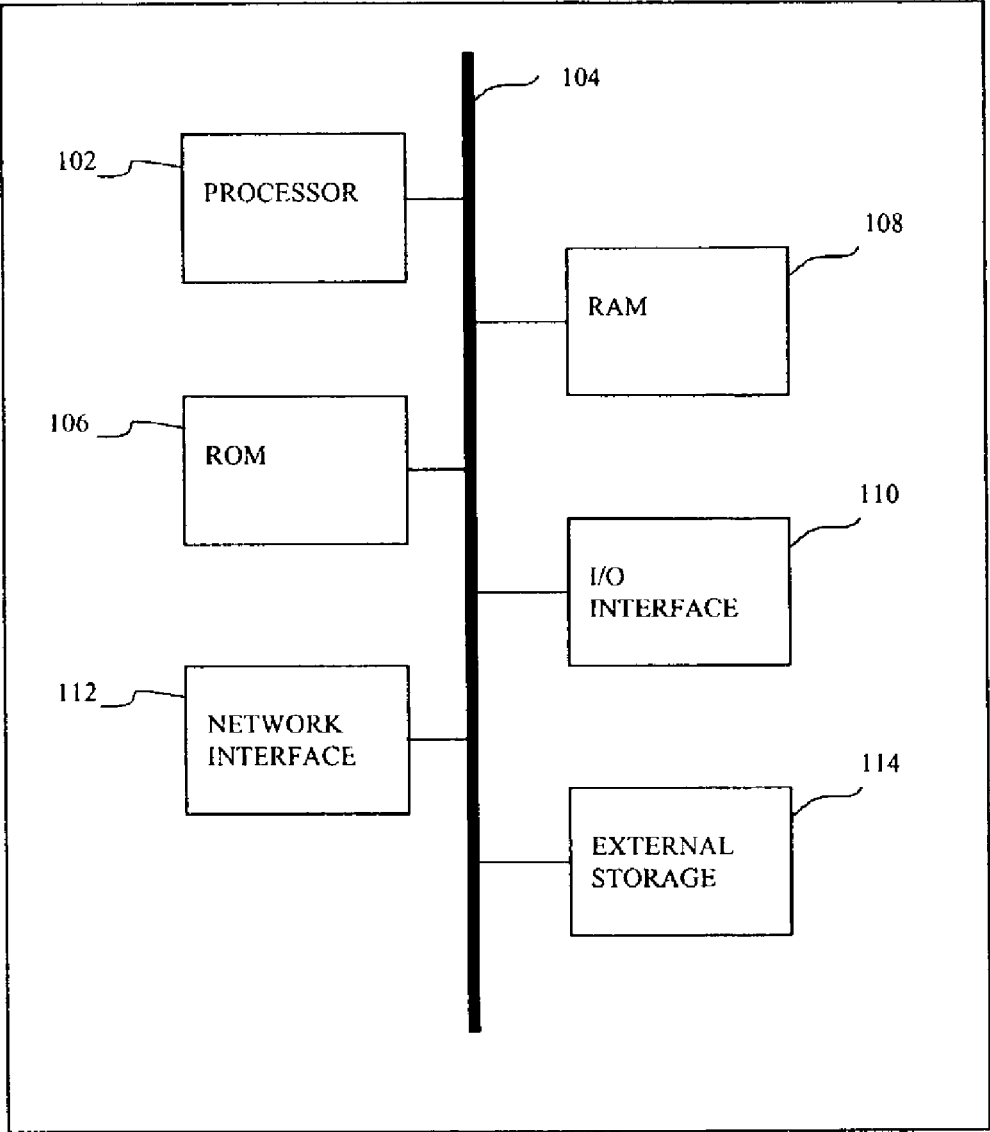
Hofer, "The Unreliabilty of Individual Physician "Report Cards" for Assessing the Costs and Quality of Care of Chronic Disease", Jun. 9, 1999, JAMA, vol. 281, No. 22:2098-2105.

Roemer et al., "Improved Diagnostic and Prognostic Assessments Using Health Management Information Fusion", 2001 IEEE Autotestcon Proceedings, IEEE Systems Readiness Technology Conference, Autotestcon 2001, Valley Forge, PA, Aug. 20-23, 2001, IEEE Systems Readiness Technology Conference, New York, NY, IEEE, US, vol. Conf. 37, Aug. 20, 2001, pp. 365-377.

Nahm et al., "Mutually Beneficial Integration of Data Mining and Information Extraction", Proceedings AAAI, National Conference on Artificial Intelligence, Jul. 30, 2000, pp. 627-632.

Rao et al., "Data Mining for Disease Management: Adding Value to Patient Records", Electromedia, Online, vol. 68, 2000, pp. 63-67.

\* cited by examiner



100 ↗

**FIG. 1**

### Patient Medical Record

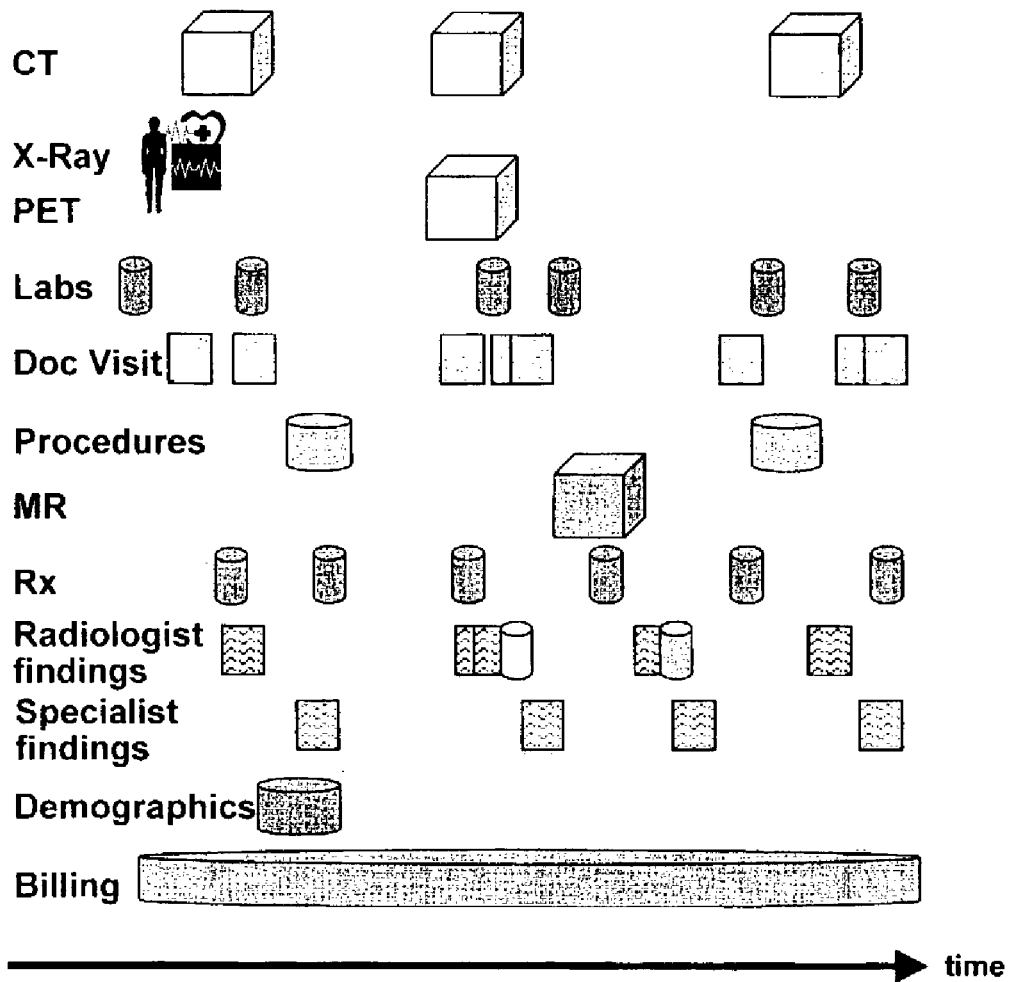


FIG. 2

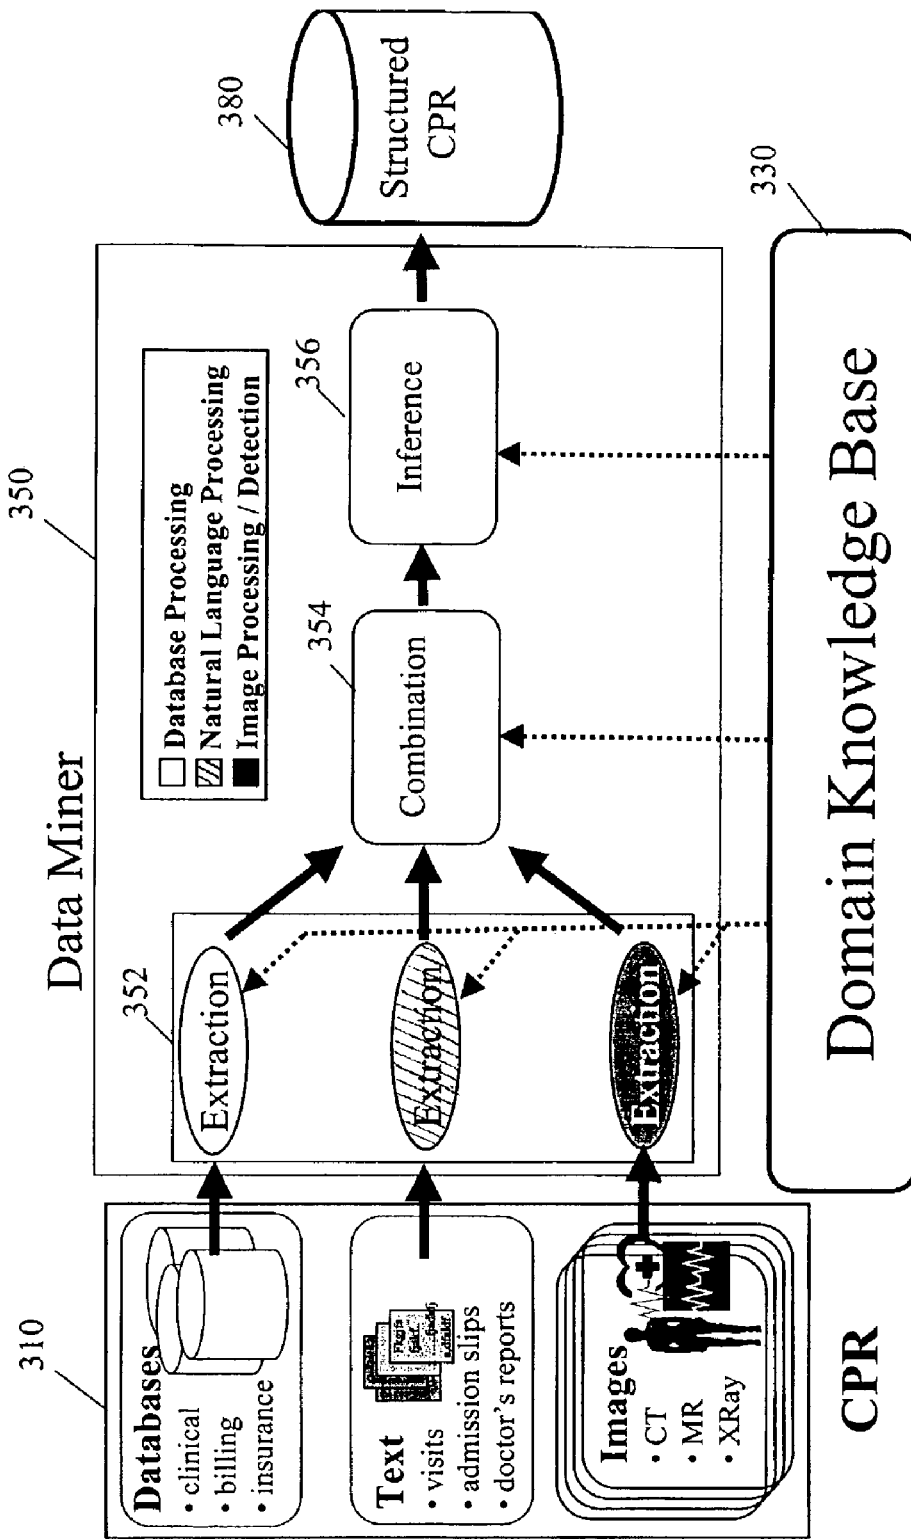


FIG 3

1

**PATIENT DATA MINING****CROSS REFERENCE TO RELATED APPLICATIONS**

This application claims the benefit of U.S. Provisional Application Serial No. 60/335,542, filed on Nov. 2, 2001, which is incorporated by reference herein in its entirety.

**FIELD OF THE INVENTION**

The present invention relates to data mining, and more particularly, to systems and methods for mining high-quality structured clinical information from patient medical records.

**BACKGROUND OF THE INVENTION**

Health care providers accumulate vast stores of clinical information. However, efforts to mine clinical information have not proven to be successful. In general, data mining is a process to determine useful patterns or relationships in data stored in a data repository. Typically, data mining involves analyzing very large quantities of information to discover trends hidden in the data.

Clinical information maintained by health care organizations is usually unstructured. Therefore, it is difficult to mine using conventional methods. Moreover, since clinical information is collected to treat patients, as opposed, for example, for use in clinical trials, it may contain missing, incorrect, and inconsistent data. Often key outcomes and variables are simply not recorded.

While many health care providers maintain billing information in a relatively structured format, this type of information is limited by insurance company requirements. That is, billing information generally only captures information needed to process medical claims, and more importantly reflects the "billing view" of the patient, i.e., coding the bill for maximum reimbursement. As a result, billing information often contains inaccurate and missing data, from a clinical point of view. Furthermore, studies show that billing codes are incorrect in a surprisingly high fraction of patients (often 10% to 20%).

Given that mining clinical information could lead to insights that otherwise would be difficult or impossible to obtain, it would be desirable and highly advantageous to provide techniques for mining structured high-quality clinical information.

**SUMMARY OF THE INVENTION**

The present invention provides a data mining framework for mining high-quality structured clinical information.

In various embodiments of the present invention, systems and methods are provided for mining information from patient records. A plurality of data sources are accessed. At least some of the data sources can be unstructured. The system includes a domain knowledge base including domain-specific criteria for mining the data sources. A data miner is configured to mine the data sources using the domain-specific criteria, to create structured clinical information.

Preferably, the data miner includes an extraction component for extracting information from the data sources to create a set of probabilistic assertions, a combination component for combining the set of probabilistic assertions to create one or more unified probabilistic assertion, and an inference component for inferring patient states from the one or more unified probabilistic assertion.

2

The extraction component may employ domain-specific criteria to extract information from the data sources. Likewise, the combination component may use domain-specific criteria to combine the probabilistic assertions, and the inference component may use domain-specific criteria to infer patient states. The patient state is simply a collection of variables that one may care about relating to the patient, for example, conditions and diagnoses.

The extraction component may be configured to extract key phrases from free text treatment notes. Other natural language processing/natural language understanding methods may also be used instead of, or in conjunction with, phrase extraction to extract information from free text.

Data sources may include one or more of medical information, financial information, and demographic information. The medical information may include one or more of free text information, medical image information, laboratory information, prescription information, and waveform information.

Probability values may be assigned to the probabilistic assertions. The structured clinical information may include probability information relating to the stored information. The structured clinical information may be stored in a data warehouse. The structured clinical information may include corrected information, including corrected ICD-9 diagnosis codes. (The International Classification of Diseases, Ninth Revision, Clinical Modification (ICD-9-CM) is based on the World Health Organization's Ninth Revision, International Classification of Diseases (ICD-9). ICD-9-CM is the official system of assigning codes to diagnosis and procedures associated with hospital utilization in the United States. The Tenth Revision (ICD-10) has recently been released and differs from the Ninth Revision (ICD-9); it is expected to be implemented soon).

The system may be run at arbitrary intervals, periodic intervals, or in online mode. When run at intervals, the data sources are mined when the system is run. In online mode, the data sources may be continuously mined.

The domain-specific criteria for mining the data sources may include institution-specific domain knowledge. For example, this may include information about the data available at a particular hospital, document structures at a hospital, policies of a hospital, guidelines of a hospital, and any variations of a hospital.

The domain-specific criteria may also include disease-specific domain knowledge. For example, the disease-specific domain knowledge may include various factors that influence risk of a disease, disease progression information, complications information, outcomes and variables related to a disease, measurements related to a disease, and policies and guidelines established by medical bodies.

Furthermore, a repository interface may be used to access at least some of the information contained in the data source used by the data miner. This repository interface may be a configurable data interface. The configurable data interface may vary depending on which hospital is under consideration.

The data source may include structured and unstructured information. Structured information may be converted into standardized units, where appropriate. Unstructured information may include ASCII text strings, image information in DICOM (Digital Imaging and Communication in Medicine) format, and text documents partitioned based on domain knowledge.

In various embodiments of the present invention, the data miner may be run using the Internet. The created structured clinical information may also be accessed using the Internet.

In various embodiments of the present invention, the data miner may be run as a service. For example, several hospitals may participate in the service to have their patient information mined, and this information may be stored in a data warehouse maintained by the service provider. The service may be performed by a third party service provider (i.e., an entity not associated with the hospitals).

These and other aspects, features and advantages of the present invention will become apparent from the following detailed description of preferred embodiments, which is to be read in connection with the accompanying drawings.

#### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a computer processing system to which the present invention may be applied according to an embodiment of the present invention;

FIG. 2 shows an exemplary computerized patient record (CPR); and

FIG. 3 shows an exemplary data mining framework for mining high-quality structured clinical information.

#### DESCRIPTION OF PREFERRED EMBODIMENTS

To facilitate a clear understanding of the present invention, illustrative examples are provided herein which describe certain aspects of the invention. However, it is to be appreciated that these illustrations are not meant to limit the scope of the invention, and are provided herein to illustrate certain concepts associated with the invention.

It is also to be understood that the present invention may be implemented in various forms of hardware, software, firmware, special purpose processors, or a combination thereof. Preferably, the present invention is implemented in software as a program tangibly embodied on a program storage device. The program may be uploaded to, and executed by, a machine comprising any suitable architecture.

Preferably, the machine is implemented on a computer platform having hardware such as one or more central processing units (CPU), a random access memory (RAM), and input/output (I/O) interface(s). The computer platform also includes an operating system and microinstruction code. The various processes and functions described herein may either be part of the microinstruction code or part of the program (or combination thereof) which is executed via the operating system. In addition, various other peripheral devices may be connected to the computer platform such as an additional data storage device and a printing device.

It is to be understood that, because some of the constituent system components and method steps depicted in the accompanying figures are preferably implemented in software, the actual connections between the system components (or the process steps) may differ depending upon the manner in which the present invention is programmed.

FIG. 1 is a block diagram of a computer processing system 100 to which the present invention may be applied according to an embodiment of the present invention. The system 100 includes at least one processor (hereinafter processor) 102 operatively coupled to other components via a system bus 104. A read-only memory (ROM) 106, a random access memory (RAM) 108, an I/O interface 110, a network interface 112, and external storage 114 are operatively coupled to the system bus 104. Various peripheral devices such as, for example, a display device, a disk storage device (e.g., a magnetic or optical disk storage device), a keyboard, and a mouse,

may be operatively coupled to the system bus 104 by the I/O interface 110 or the network interface 112.

The computer system 100 may be a standalone system or be linked to a network via the network interface 112. The network interface 112 may be a hard-wired interface. However, in various exemplary embodiments, the network interface 112 can include any device suitable to transmit information to and from another device, such as a universal asynchronous receiver/transmitter (UART), a parallel digital interface, a software interface or any combination of known or later developed software and hardware. The network interface may be linked to various types of networks, including a local area network (LAN), a wide area network (WAN), an intranet, a virtual private network (VPN), and the Internet.

The external storage 114 may be implemented using a database management system (DBMS) managed by the processor 102 and residing on a memory such as a hard disk. However, it should be appreciated that the external storage 114 may be implemented on one or more additional computer systems. For example, the external storage 114 may include a data warehouse system residing on a separate computer system.

Those skilled in the art will appreciate that other alternative computing environments may be used without departing from the spirit and scope of the present invention.

Increasingly, health care providers are employing automated techniques for information storage and retrieval. The use of a computerized patient record (CPR) to maintain patient information is one such example. As shown in FIG. 2, an exemplary CPR (200) includes information that is collected over the course of a patient's treatment. This information may include, for example, computed tomography (CT) images, X-ray images, laboratory test results, doctor progress notes, details about medical procedures, prescription drug information, radiological reports, other specialist reports, demographic information, and billing (financial) information.

A CPR typically includes a plurality of data sources, each of which typically reflect a different aspect of a patient's care. Structured data sources, such as financial, laboratory, and pharmacy databases, generally maintain patient information in database tables. Information may also be stored in unstructured data sources, such as, for example, free text, images, and waveforms. Often, key clinical findings are only stored within physician reports.

FIG. 3 illustrates an exemplary data mining system for mining high-quality structured clinical information. The data mining system includes a data miner (350) that mines information from a CPR (310) using domain-specific knowledge contained in a knowledge base (330). The data miner (350) includes components for extracting information from the CPR (352), combining all available evidence in a principled fashion over time (354), and drawing inferences, from this combination process (356). The mined information may be stored in a structured CPR (380).

The extraction component (352) deals with gleaning small pieces of information from each data source regarding a patient, which are represented as probabilistic assertions about the patient at a particular time. These probabilistic assertions are called elements. The combination component (354) combines all the elements that refer to the same variable at the same time period to form one unified probabilistic assertion regarding that variable. These unified probabilistic assertions are called factoids. The inference component (356) deals with the combination of these factoids, at the same point in time and/or at different points in time, to produce a coher-



ent and concise picture of the progression of the patient's state over time. This progression of the patient's state is called a state sequence.

The present invention can build an individual model of the state of a patient. The patient state is simply a collection of variables that one may care about relating to the patient. The information of interest may include a state sequence, i.e., the value of the patient state at different points in time during the patient's treatment.

Advantageously, the architecture depicted in FIG. 3 supports plug-in modules wherein the system can be easily expanded for new data sources, diseases, and hospitals. New element extraction algorithms, element combining algorithms, and inference algorithms can be used to augment or replace existing algorithms.

Each of the above components uses detailed knowledge regarding the domain of interest, such as, for example, a disease of interest. This domain knowledge base (330) can come in two forms. It can be encoded as an input to the system, or as programs that produce information that can be understood by the system. The part of the domain knowledge base (330) that is input to the present form of the system may also be learned from data.

Domain-specific knowledge for mining the data sources may include institution-specific domain knowledge. For example, this may include information about the data available at a particular hospital, document structures at a hospital, policies of a hospital, guidelines of a hospital, and any variations of a hospital.

The domain-specific knowledge may also include disease-specific domain knowledge. For example, the disease-specific domain knowledge may include various factors that influence risk of a disease, disease progression information, complications information, outcomes-and variables related to a disease, measurements related to a disease, and policies and guidelines established by medical bodies.

As mentioned, the extraction component (352) takes information from the CPR (310) to produce probabilistic assertions (elements) about the patient that are relevant to an instant in time or time period. This process is carried out with the guidance of the domain knowledge that is contained in the domain knowledge base (330). The domain knowledge required for extraction is generally specific to each source.

Extraction from a text source may be carried out by phrase spotting, which requires a list of rules that specify the phrases of interest and the inferences that can be drawn therefrom. For example, if there is a statement in a doctor's note with the words, "There is evidence of metastatic cancer in the liver," then, in order to infer from this sentence that the patient has cancer, a rule is needed that directs the system to look for the phrase "metastatic cancer," and, if it is found, to assert that the patient has cancer with a high degree of confidence (which, in the present embodiment, translates to generate an element with name "Cancer", value "True" and confidence 0.9).

The data sources include structured and unstructured information. Structured information may be converted into standardized units, where appropriate. Unstructured information may include ASCII text strings, image information in DICOM (Digital Imaging and Communication in Medicine) format, and text documents partitioned based on domain knowledge. Information that is likely to be incorrect or missing may be noted, so that action may be taken. For example, the mined information may include corrected information, including corrected ICD-9 diagnosis codes.

Extraction from a database source may be carried out by querying a table in the source, in which case, the domain knowledge needs to encode what information is present in

which fields in the database. On the other hand, the extraction process may involve computing a complicated function of the information contained in the database, in which case, the domain knowledge may be provided in the form of a program that performs this computation whose output may be fed to the rest of the system.

Extraction from images, waveforms, etc., may be carried out by image processing or feature extraction programs that are provided to the system.

Combination includes the process of producing a unified view of each variable at a given point in time from potentially conflicting assertions from the same/different sources. In various embodiments of the present invention, this is performed using domain knowledge regarding the statistics of the variables represented by the elements ("prior probabilities").

Inference is the process of taking all the factoids that are available about a patient and producing a composite view of the patient's progress through disease states, treatment protocols, laboratory tests, etc. Essentially, a patient's current state can be influenced by a previous state and any new composite observations.

The domain knowledge required for this process may be a statistical model that describes the general pattern of the evolution of the disease of interest across the entire patient population and the relationships between the patient's disease and the variables that may be observed (lab test results, doctor's notes, etc.). A summary of the patient may be produced that is believed to be the most consistent with the information contained in the factoids, and the domain knowledge.

For instance, if observations seem to state that a cancer patient is receiving chemotherapy while he or she does not have cancerous growth, whereas the domain knowledge states that chemotherapy is given only when the patient has cancer, then the system may decide either: (1) the patient does not have cancer and is not receiving chemotherapy (that is, the observation is probably incorrect), or (2) the patient has cancer and is receiving chemotherapy (the initial inference—that the patient does not have cancer—is incorrect); depending on which of these propositions is more likely given all the other information. Actually, both (1) and (2) may be concluded, but with different probabilities.

As another example, consider the situation where a statement such as "The patient has metastatic cancer" is found in a doctor's note, and it is concluded from that statement that  $\langle \text{cancer}=\text{True} \text{ (probability}=0.9) \rangle$ . (Note that this is equivalent to asserting that  $\langle \text{cancer}=\text{True} \text{ (probability}=0.9) \rangle$ ,  $\langle \text{cancer}=\text{unknown} \text{ (probability}=0.1) \rangle$ ).

Now, further assume that there is a base probability of cancer  $\langle \text{cancer}=\text{True} \text{ (probability}=0.35) \rangle$ ,  $\langle \text{cancer}=\text{False} \text{ (probability}=0.65) \rangle$  (e.g., 35% of patients have cancer). Then, we could combine this assertion with the base probability of cancer to obtain, for example, the assertion  $\langle \text{cancer}=\text{True} \text{ (probability}=0.93) \rangle$ ,  $\langle \text{cancer}=\text{False} \text{ (probability}=0.07) \rangle$ .

Similarly, assume conflicting evidence indicated the following:

1.  $\langle \text{cancer}=\text{True} \text{ (probability}=0.9) \rangle$ ,  $\langle \text{cancer}=\text{unknown} \text{ (probability}=0.1) \rangle$
2.  $\langle \text{cancer}=\text{False} \text{ (probability}=0.7) \rangle$ ,  $\langle \text{cancer}=\text{unknown} \text{ (probability}=0.3) \rangle$
3.  $\langle \text{cancer}=\text{True} \text{ (probability}=0.1) \rangle$ ,  $\langle \text{cancer}=\text{unknown} \text{ (probability}=0.9) \rangle$  and
4.  $\langle \text{cancer}=\text{False} \text{ (probability}=0.4) \rangle$ ,  $\langle \text{cancer}=\text{unknown} \text{ (probability}=0.6) \rangle$ .

In this case, we might combine these elements with the base probability of cancer  $\langle \text{cancer}=\text{True} \text{ (probability}=0.35) \rangle$ ,

cancer=False (probability=0.65)>to conclude, for example, that <cancer=True (prob=0.67), cancer=False (prob=0.33)>.

It should be appreciated the present invention typically must access numerous data sources, and deal with missing, incorrect, and/or inconsistent information. As an example, consider that, in determining whether a patient has diabetes, the following information might have to be extracted:

(a) ICD-9 billing codes for secondary diagnoses associated with diabetes;

(b) drugs administered to the patient that are associated with the treatment of diabetes (e.g., insulin);

(c) patient's lab values that are diagnostic of diabetes (e.g., two successive blood sugar readings over 250 mg/d);

(d) doctor mentions that the patient is a diabetic in the H&P (history & physical) or discharge note (free text); and

(e) patient procedures (e.g., foot exam) associated with being a diabetic.

As can be seen, there are multiple independent sources of information, observations from which can support (with varying degrees of certainty) that the patient is diabetic (or more generally has some disease/condition). Not all of them may be present, and in fact, in some cases, they may contradict each other. Probabilistic observations can be derived, with varying degrees of confidence. Then these observations (e.g., about the billing codes, the drugs, the lab tests, etc.) may be probabilistically combined to come up with a final probability of diabetes. Note that there may be information in the patient record that contradicts diabetes. For instance, the patient is has some stressful episode (e.g., an operation) and his blood sugar does not go up.

It should be appreciated that the above examples are presented for illustrative purposes only and are not meant to be limiting. The actual manner in which elements are combined depends on the particular domain under consideration as well as the needs of the users of the system. Further, it should be appreciated that while the above discussion refers to a patient-centered approach, actual implementations may be extended to handle multiple patients simultaneously. Additionally, it should be appreciated that a learning process may be incorporated into the domain knowledge base (330) for any or all of the stages (i.e., extraction, combination, inference) without departing from the spirit and scope of the present invention.

The system may be run at arbitrary intervals, periodic intervals, or in online mode. When run at intervals, the data sources are mined when the system is run. In online mode, the data sources may be continuously mined.

The data miner may be run using the Internet. The created structured clinical information may also be accessed using the Internet.

Additionally, the data miner may be run as a service. For example, several hospitals may participate in the service to have their patient information mined, and this information may be stored in a data warehouse owned by the service provider. The service may be performed by a third party service provider (i.e., an entity not associated with the hospitals).

Once the structured CPR (380) is populated with patient information, it will be in a form where it is conducive for answering several questions regarding individual patients, and about different cross-sections of patients.

The following describes REMIND (Reliable Extraction and Meaningful Inference from Non-structured Data), an innovative data mining system developed by Siemens Corporate Research (SCR), a subsidiary of Siemens Corporation. REMIND is based upon an embodiment of the present invention.

Initially, an analogy is provided that describes the spirit in which REMIND performs inferences.

A French medical student who has some knowledge about cancer is provided with cancer patient CPR's. The CPR's contain transcribed English dictations and pharmacy data. The student's task is to classify which patients have had a recurrence, and if they have, determine when it occurred. Unfortunately his English is poor, though he does know some key medical words and a few of the drug names. However, he cannot rely purely on the presence of some key words, such as metastases, in the dictation, because he knows that physicians often make negative statements ("Patient is free of evidence of metastases."). How might the student best carry out his task?

The student can collect all relevant evidence from the CPR—without trusting any single piece of evidence—and combine it to reconcile any disparities. He can use his knowledge about the treatment of cancer—for instance, on noting that a patient had a liver resection, the student can conclude that the patient (probably) previously had a recurrence.

#### Problem Definition

Let  $S$  be a continuous time random process taking values in  $\Sigma$  that represents the state of the system. Let  $T=\{t_1, t_2, \dots, t_n\}$ , where  $t_i < t_{i+1}$ , be the  $n$  "times of interest" when  $S$  has to be inferred. Let  $S_t$  refer to the sample of  $S$  at time  $t_i \in T$ . Let  $V$  be the set of variables that depend upon  $S$ . Let  $O$  be set of all (probabilistic) observations for all variables,  $v \in V$ . Let  $O_i$  be the set of all observations "assigned" to  $t_i \in T$ ; i.e., all observations about variables,  $v \in V$ , that are relevant for this time-step  $t_i$ . Similarly, let

$$\text{seq}_{MAP} = \arg \max_{\text{seq}} P[\text{seq} | O]$$

$O'_j(v)$  be the  $j$ -th observation for variable  $v$  assigned to  $t_j$ . Let  $\text{seq} = \langle S_1, S_2, \dots, S_n \rangle$  be a random variable in  $\Sigma^n$ ; i.e., each realization of  $\text{seq}$  is a state sequence across  $T$ . GOAL: Estimate the most likely state sequence,  $\text{seq}_{MAP}$ , (the maximum a posteriori estimate of  $\text{seq}$ ) given  $O$ .

REMIND extracts information,  $o_i$ , from every data source in a uniform format called probabilistic observations. Each  $o_i$  is drawn entirely from a single piece of information in a data source (e.g., from a phrase in a sentence, or a row in a database table), and hence is assumed to be inherently undependable. The observation {"Recurrent", "12/17/01", <T=0.1, F=0.0>}, states that the Boolean variable "Recurrent" has an associated distribution over all possible values that can be taken by "Recurrent". The probabilities do not have to add up to 1.0; any remainder (here 0.9) is assigned to unknown, and is smoothed over T/F, based upon the (time-dependent) a priori distribution.

Extraction from Structured data: REMIND communicates with all databases via JDBC, Java's built-in interface to relational databases. Executing a query (e.g., retrieve drug administered) is expressed as a probabilistic observation.

Extraction from Free Text: REMIND strips document headers/footers, and tokenizes free text. Information from the token stream is extracted via phrase spotting, an easy-to-implement method from computational linguistics. Phrase spotting is about as simple as it sounds. A phrase-spotting rule is applied within a single sentence. The rule:

[metastasis & malignant]  $\Rightarrow$  {"Recurrent", <T=0.5>} states that if the 2 words (actually aliases) in the rule are found in a sentence, a probabilistic observation about recurrence should be generated. REMIND also has compound rules to detect "negation" and "imprecision", which modify the probabilities in existing observations.

The primary focus of our interest is estimating what happened to the patient across T, the duration of interest. The estimation of the MAP state sequence can be done in two steps, the first of which is combination of observations at a fixed point in time and the second is the propagation of these inferences across time.

Each (smoothed)  $o_t$  is in the form of an a posteriori probability of a variable given the small context that it is extracted from. All observations,  $O'_i(v)$ , about a variable for a single time  $t_i$  are combined into one assertion in a straightforward manner by using Bayes' theorem:

$$P[v_i | O'_1(v_i), \dots, O'_k(v_i)] \propto$$

$$P[v_i] \cdot \prod_{j=1}^k P[O'_j(v_i) | v_i] \propto \frac{\prod_{j=1}^k P[v_i | O'_j(v_i)]}{P[v_i]^{k-1}}$$

At every  $t_i \in T$ , the relationships among S, and V are modeled using a Bayesian Network. Because the state process is modeled as being Markov and the state as being causative (directly or indirectly) of all the variables that we observe, we have the following equation:

$$P[seq | O] \propto P[S_0] \cdot \prod_{i=2}^n P[S_i | S_{i-1}] \cdot \prod_{i=1}^n P[O_i | S_i] \propto \prod_{i=2}^n \frac{P[S_i | S_{i-1}]}{P[S_i]} \cdot \prod_{i=1}^n P[S_i | O_i]$$

This equation connects the a posteriori probability of seq (any sequence of samples of the state process across time) given all observations, to  $P(S_i | O_i)$ , the temporally local a posteriori probability of the state given the observations for each time instant. Essentially, we string together the temporally local Bayesian Networks by modeling each state sample,  $S_{t_i}$ , as the cause of the next sample,  $S_{t_{i+1}}$ .

Although illustrative embodiments of the present invention have been described herein with reference to the accompanying drawings, it is to be understood that the invention is not limited to those precise embodiments, and that various other changes and modifications may be affected therein by one skilled in the art without departing from the scope or spirit of the invention.

What is claimed is:

1. A system for producing structured clinical information from patient records, comprising:

a computerized patient record comprising at least one data source storing patient information, at least some of the patient information being unstructured, the patient information including clinical information related to a plurality of variables representing a state of the patient; a memory comprising a domain knowledge base containing domain-specific criteria relating the clinical information to values of the variables; and

a processor configured to apply a data miner for extracting the clinical information from the data source using the domain-specific criteria, the data miner operable to create structured clinical information from the extracted clinical information including at least some of the

unstructured patient information, the structured clinical information including the variables representing a state of the patient as a summary, the data miner operable to resolve discrepancies of different clinical information indicating different values of one of the variables.

2. The system of claim 1, wherein the data miner comprises:

an extraction component for extracting information from the data sources and outputting a set of probabilistic assertions;

a combination component for combining the set of probabilistic assertions into one or more unified probabilistic assertions; and

an inference component for inferring patient states from the one or more unified probabilistic assertions.

3. The system of claim 2, wherein the extraction component extracts from the data sources as a function of the domain-specific criteria.

4. The system of claim 2, wherein the combination component combines from the probabilistic assertions as a function of the domain-specific criteria.

5. The system of claim 2, wherein the inference component infers the patient states as a function of the domain-specific criteria.

6. The system of claim 1, wherein the data sources include one or more of: medical information, financial information, demographic information, billing information or combinations thereof.

7. The system of claim 6, wherein the medical information includes one or more of: free text information, medical image information, laboratory information, prescription drug information, waveform information or combinations thereof.

8. The system of step 1, wherein the data miner is run at arbitrary intervals.

9. The system of claim 1, wherein the data miner is run at periodic intervals.

10. The system of step 1, wherein the data miner is run in online mode.

11. The system of claim 2, wherein the extraction component extracts key phrases from free text treatment notes.

12. The system of claim 2, wherein probability values are assigned to the probabilistic assertions.

13. The system of claim 1, wherein the created structured clinical information is stored in a database.

14. The system of claim 1, wherein the created structured clinical information includes probability information.

15. The system of claim 2, wherein the inference component uses a statistical model that describes a pattern of evolution of a disease across a patient population and the relationship between a patient's disease and observed variables.

16. The system of claim 15, wherein the inference component draws a plurality of inferences, each with an assigned probability.

17. The system of claim 1, wherein the domain-specific criteria includes institution-specific domain knowledge.

18. The system of claim 17, wherein the institution-specific domain knowledge relates to one or more of: data at a hospital, document structures at a hospital, policies of a hospital, guidelines of a hospital, variations at a hospital or combinations thereof.

19. The system of claim 1, wherein the domain-specific criteria includes disease-specific domain knowledge.

20. The system of claim 19, wherein the disease-specific domain knowledge includes one or more of: factors that influence risk of a disease, disease progression information, complications information, outcomes related to the disease, variables related to a disease, measurements related to a disease,

## 11

policies established by medical bodies, guidelines established by medical bodies or combinations thereof.

21. The system of claim 1, wherein at least some of the information contained in the data source is accessible with a repository interface.

22. The system of claim 21, wherein the repository interface is a configurable data interface.

23. The system of claim 22, wherein the configurable data interface varies depending on hospital.

24. The system of claim 1, wherein the data sources include structured information, the created structured clinical information being, in part, responsive to the structured information.

25. The system of claim 24, wherein the structured information is converted into standardized units.

26. The system of claim 1, wherein the unstructured information includes one or more of: ASCII text strings, image information in DICOM format, text documents or combinations thereof partitioned based on domain knowledge.

27. The system of claim 1, wherein the data miner is run using the Internet.

28. The system of claim 1, wherein the created structured clinical information is accessed using the Internet.

29. The system of claim 1, wherein the data miner is run as a service.

30. The system of claim 29, wherein the service is performed by a third party service provider.

31. The system of claim 2, wherein the inferred patient states include diagnoses.

32. The system of claim 1, wherein the created structured clinical information includes corrected information.

33. The system of claim 1, wherein the data miner comprises an extraction component for extracting information from the data sources and outputting a set of probabilistic assertion.

34. The system of claim 33, wherein the data miner comprises a component for combining the set of probabilistic assertions into one or more unified probabilistic assertions and for inferring patient states from the one or more unified probabilistic assertions.

35. The system of claim 34 wherein the component for combining and for inferring is operable pursuant to Bayes Theorem.

36. The system of claim 2 wherein the inference and combination components are operable pursuant to Bayes Theorem.

37. The system of claim 1 wherein the at least one data source comprises a health care provider patient record.

38. The system of claim 13 wherein the database comprises a healthcare provider database.

39. The system of claim 11, wherein the extraction component extracts the key phrases using phrase spotting.

40. The system of claim 1, wherein the domain-specific criteria comprises criteria learned from data.

41. A method for producing structured clinical information from patient records, comprising the steps of:

providing a first computerized patient record comprising at least one data source containing patient information, at least some of the patient information being unstructured and at least some of the patient information being structured;

providing a memory storing a domain knowledge base containing domain-specific criteria;

extracting clinical information from the at least one data source by data mining using the domain-specific criteria; and

## 12

creating, using a computer, structured clinical information as a function of the extracting, the structured clinical information being a summary of the unstructured and structured patient information.

42. The method of claim 41, wherein extracting the clinical information from the data sources comprises:

creating a set of probabilistic assertions;

combining the set of probabilistic assertions into one or more unified probabilistic assertions; and

inferring patient states from the one or more unified probabilistic assertions.

43. The method of claim 42, wherein extracting information from the data sources includes extracting the extracted information from the data sources as a function of the domain-specific criteria.

44. The method of claim 42, wherein combining the set of probabilistic assertions includes combining the probabilistic assertions as a function of the domain-specific criteria.

45. The method of claim 42, wherein inferring the patient states includes inferring the patient states as a function of the domain-specific criteria.

46. The method of claim 41, wherein the data sources include one or more of: medical information, financial information, demographic information, billing information or combinations thereof.

47. The method of claim 46, wherein the medical information includes one or more of: free text information, medical image information, laboratory information, prescription drug information, waveform information or combinations thereof.

48. The method of claim 42, wherein probability values are assigned to the probabilistic assertions.

49. The method of claim 41, wherein the created structured clinical information is stored in a database.

50. The method of claim 41, wherein the created structured clinical information includes probability information.

51. The method of claim 41, wherein the domain-specific criteria for mining the data sources includes institution-specific domain knowledge.

52. The method of claim 51, wherein the institution-specific domain knowledge relates to one or more of: data at a hospital, document structures at a hospital, policies of a hospital, guidelines of hospital, variations at a hospital or combinations thereof.

53. The method of claim 41, wherein the domain-specific criteria includes disease-specific domain knowledge.

54. The method of claim 42, combining and inferring comprises using Bayes Theorem.

55. The method of claim 41 wherein providing at least one data source comprises providing a health care provider patient record.

56. The method of claim 49 wherein the database comprises a healthcare provider database.

57. The method of claim 51, wherein the disease-specific domain knowledge includes one or more of: factors that influence risk of a disease, disease progression information, complications information, outcomes related to the disease, variables related to a the disease, measurements related to the disease, policies established by medical bodies, guidelines established by medical bodies or combinations thereof.

58. The method of claim 41, wherein the data source includes structured information.

59. The method of claim 58, wherein the structured information is converted into standardized units.

60. The method of claim 41, wherein the unstructured information includes one or more of: ASCII text strings,

**13**

image information in DICOM format, text documents or combinations thereof partitioned based on domain knowledge.

61. The method of claim 41, wherein the method is performed using the Internet.

62. The method of claim 41, wherein the method is performed by a third party service provider.

63. The method of claim 42, wherein the inferred patient states include diagnoses.

64. The method of claim 41, wherein the created structured clinical information includes corrected information.

65. A program storage device readable by a machine, tangibly embodying a program of instructions when executed on the machine to perform a method for producing structured clinical information from patient records, the method comprising:

**14**

providing at least one data source containing patient information, at least some of the patient information being unstructured;

providing a domain knowledge base containing domain-specific criteria;

extracting clinical information from the data sources using the domain-specific criteria, wherein extracting the clinical information from the data sources comprises: creating a set of probabilistic assertions; combining the set of probabilistic assertions into one or more unified probabilistic assertions; and inferring patient states from the one or more unified probabilistic assertions; and creating structured clinical information as a function of the extracting.

\* \* \* \* \*