



- (51) International Patent Classification:
G06Q 50/00 (2012.01)
- (21) International Application Number:
PCT/US2017/059025
- (22) International Filing Date:
30 October 2017 (30.10.2017)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
15/337,923 28 October 2016 (28.10.2016) US
- (71) Applicant: INTELLIGENT MEDICAL OBJECTS, INC. [US/US]; 60 Revere Drive, Suite 400, Northbrook, Illinois 60062 (US).
- (72) Inventors: NAEYMI-RAD, Frank; 1310 St. Williams Drive, Libertyville, Illinois 60048 (US). CARDWELL, Matthew; 315 N. East Ave, Oak Park, Illinois 60302 (US). DECARO, Michael; 1706 E. Mohawk Lane, Mount Prospect, Illinois 60056 (US). PAUL, Erina; 1400 W Lincoln Hwy Apt B6, DeKalb, Illinois 60115 (US). WANG, Yiqing; 1400 Twombly Road Apt 1735, DeKalb, Illinois 60115 (US). THOMPSON, James; 1508 Triple Crown Court, St. Charles, Illinois 60174 (US). KANTER, An-

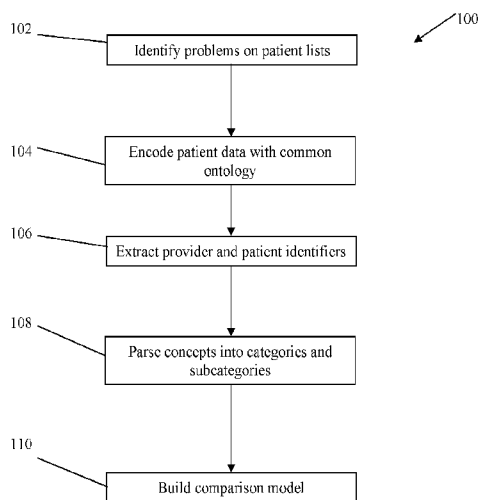
drew; 1755 Lake Cook Rd, Apt. 317, Highland Park, Illinois 60035 (US).

(74) Agent: LINZER, John R; Quarles & Brady LLP (CHI), Attn: IP Docket, 411 E. Wisconsin Ave., Suite 2350, Milwaukee, Wisconsin 53202-4426 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,

(54) Title: EXTRACTING PATIENT DATA TO PROVIDE PROVIDER AND PATIENT DATA SIMILARITY SCORING



(57) Abstract: A system and method for extracting data from an electronic health record to provide provider and patient data similarity scoring includes: encoding a problem list for a plurality of patients with concepts from a common electronic health record ontology. In one aspect, the patients have electronic health records maintained by a plurality of providers. The system and method then may parse the concepts into a plurality of clusters or categories and determining, for each of the providers, a total number of patients that have at least one problem in a cluster or category or determining, for each patient, which of the plurality of clusters or categories correspond to at least one concept encoded in the patient's problem list. The system and method then may calculate for each pair of providers or patients, a distance between the providers or patients.

FIG. 1

WO 2018/081703 A1

TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

EXTRACTING PATIENT DATA TO PROVIDE PROVIDER AND PATIENT DATA SIMILARITY SCORING

BACKGROUND

1. Field of the Invention

[0001] The present application is directed to electronic health record or other electronic patient data repository tools, including analytical tools for effectively using data encoded within electronic health records or other data repositories.

2. Description of the Related Art

[0002] Electronic medical ontologies, also known as medical classification codes, are necessary with the implementation and proliferation of electronic medical records. Various ontologies have been developed for various reasons, including administrative code sets that may be designed to support administrative functions of healthcare, such as reimbursement and other secondary data aggregation; clinical code sets that encode specific clinical entities involved in clinical work flow and allow for meaningful electronic exchange and aggregation of clinical data for better patient care; and reference terminology code sets that may be considered a "concept-based, controlled medical terminology" to maintain a common reference point in the healthcare industry. Reference terminologies also identify relationships between their concepts, e.g., relationships can be hierarchically defined, such as a parent/child relationship. Common examples of administrative code sets are the International Classification of Disease (ICD) and the Current Procedural Terminology, which is referred to via the trademark CPT. Examples of clinical code sets are the Logical Observation Identifiers Names and Codes, referred to under the trademark LOINC, and a normalized terminology for medication information, such as the terminology of the National Library of Medicine referred to under the trademark RxNorm. One example of a reference terminology is The Systematized Nomenclature of Medicine - Clinical Terms, referred to under the trademark "SNOMED CT."

[0003] One challenge with implementing an electronic medical ontology is to ensure the accuracy and completeness of recordkeeping, at the time of the patient visit or otherwise during data entry. One method of structuring and codifying the data to achieve this goal includes implementing an interface terminology that recognizes semantic meaning, mapping that interface terminology to the various other ontologies, and then relying on that interface terminology to analyze the practitioner's entries. One example of a system and method for using an interface terminology and the relevant ontology mappings may be found in the commonly-owned U.S. patent publication 2014/0122117, published May 1, 2014, the contents of which are incorporated

by reference in their entirety. In that example, the interface terminology comprises a plurality of concepts within one or more domains, and one or more descriptions (lexicals) linked to each concept, where each description reflects an alternative way to express the concept.

[0004] Separately, care providers across all disciplines see myriad patients presenting with various problems. Often, treatment of those patients is fairly straightforward, as the problems are those with which the care providers deal frequently or for which treatment plans are commonly known. At times, however, the provider may face a problem or combination of problems with an unknown or more rear treatment protocol, and it may be desirable to draw upon the knowledge and experience of other providers in addressing those problems. While the ontologies described above may be useful for recordkeeping, billing, etc., the sheer volume of codes or entries within each ontology may hinder searching and meaningful analysis. For example, within ICD-10-CM, there are almost 70,000 distinct diagnosis codes to which particular problems can be mapped, and there are an additional 87,000 distinct codes within ICD-10-PCS for mapping procedure-related data. Additionally, due to various other constraints, including geography, patient confidentiality, etc., finding the relevant information may be difficult and time consuming, if not seemingly impossible.

[0005] Still further, instead of looking outward, care providers may be able to draw upon their own internal knowledge and previous experience in addressing a problem or combination of problems. In this regard, the volume of patients seen by the provider or the amount of time elapsed between the current instance of a problem or combination of problems and the relevant previous instance may hamper the provider's recall of the relevant details.

[0006] What are needed are a system and method that preferably address one or more of these challenges.

BRIEF SUMMARY

[0007] In one aspect, a method for extracting data from an electronic health record to provide provider and patient data similarity scoring includes: encoding a problem list for a plurality of patients with concepts from a common electronic health record ontology, wherein the plurality of patients have electronic health records maintained by a plurality of providers, parsing the concepts into a plurality of clusters or categories, determining, for each of the providers, a total number of patients that have at least one problem in a cluster or category, iterating the determining step for each of the remaining clusters or categories, and calculating, for each pair of providers, a distance between the providers.

[0008] In another aspect, a method for extracting data from an electronic health record to provide provider and patient data similarity scoring includes: encoding a problem list for a plurality of patients with concepts from a common electronic health record ontology, parsing the concepts into a plurality of clusters or categories, determining, for each patient, which of the plurality of clusters or categories correspond to at least one concept encoded in the patient's problem list, and calculating, for each pair of patients, a distance between the patients.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

[0009] FIG. 1 is a flowchart depicting one example of a method for extracting data from a patient problem list, e.g., a problem list stored in or accessible by an electronic health record, to provide provider and patient data similarity scoring.

[0010] FIG. 2 is a depiction of a method of reconciling a general problem list into one or more clinical categories based on concept groupings. In this case the concept group is related to clinical specialties such as Gastroenterology or Cardiovascular. Many different concept groupings can be enabled using the methods described.

[0011] FIG. 3 is a depiction of exemplary relationships between problem list elements within a clinical category and an example of how problems can be nested together or seen in full detail.

[0012] FIG. 4 is a depiction of problem lists from different sources, illustrating differences in the way in which problem list elements are arranged and displayed.

[0013] FIG. 5 is a flowchart depicting a subroutine of the method of FIG. 1 directed to a provider-centric aspect of the method.

[0014] FIG. 6 is a flowchart depicting a subroutine of the method of FIG. 1 directed to a patient-centric aspect of the method.

[0015] FIG. 7 is a visual depiction of a plurality of providers with problem list concepts combined into clusters and arranged to indicate similarity between providers.

[0016] FIG. 8 is a visual depiction of a selected provider from among the providers of FIG. 4 and the providers most similar to the selected provider.

[0017] FIG. 9 is a visual depiction of a plurality of providers with problem list concepts combined into categories and arranged to indicate similarity between providers.

[0018] FIG. 10 is a visual depiction of a selected provider from among the providers of FIG. 6 and the providers most similar to the selected provider.

[0019] FIG. 11 is a visual depiction of a plurality of providers with problem list concepts recorded in a limited time span, combined into clusters, and arranged to indicate similarity between providers.

[0020] FIG. 12 is a visual depiction of a selected patient from among a plurality of patients with problem list concepts combined into clusters, the depiction arranged to indicate similarity between patients.

[0021] FIG. 13 is a visual depiction of a selected patient from among a plurality of patients with problem list concepts combined into categories, the depiction arranged to indicate similarity between patients.

[0022] FIG. 14 is a visual depiction of a selected patient from among a plurality of patients analyzed for exact matches to problem list concepts, the depiction arranged to indicate similarity between patients.

[0023] FIG. 15 is a visual depiction of a selected patient from among a plurality of patients analyzed, accounting equally for problem list concept matches and cluster similarity, the depiction arranged to indicate similarity between patients.

[0024] FIG. 16 is a visual depiction of a selected patient from among a plurality of patients analyzed, accounting more heavily for problem list concept matches as compared to cluster similarity, the depiction arranged to indicate similarity between patients.

[0025] FIG. 17 is a visual depiction of a selected patient from among a plurality of patients analyzed, accounting more heavily for cluster similarity as compared to problem list concept matches, the depiction arranged to indicate similarity between patients.

DETAILED DESCRIPTION

[0026] As set forth in greater detail herein, the present system and method are operable within a network of computer systems, with a plurality of computers each having a processor configured to operate electronic health record ("EHR") software accessible by one or more care providers to document patient encounters. In one aspect, each computer system operates the same EHR software. In another aspect, the computer systems may operate different EHR software packages that receive and/or store patient data in different ways. In this latter aspect, however, the various EHR software packages may interface with a common ontology such as an interface terminology in order to provide a common encoding mechanism for their respective sets of patient data.

[0027] The computer systems may be interconnected as part of a common network servicing a single entity, such as a solo practitioner or a physician's office. Alternatively, the computer systems may be separately configured to serve a plurality of related entities, such as a hospital or a group of affiliated providers. In yet another option, the computer systems may be configured to serve a plurality of unrelated entities, such as unaffiliated care providers. In each instance, patient records may be stored in a database in each computer system or in a centralized database

operatively coupled to each respective computer. Additionally, in any of these instances, the computer systems may be in communication with a centralized computer system, which may include one or more servers configured to distribute the common ontology to each of the computer systems. The centralized computer system also may receive patient-related data from each of the other computer systems.

[0028] Each EHR software package may include, as one component, a patient problem list that maintains a history of all patient problems. The problem list may distinguish between current and past problems, as well as chronic versus acute problems, in order to give providers an accurate picture of each patient's history. Each entry in the problem list may be encoded with an entry from one or more different ontologies, permitting the EHR to interact with other software applications. For example, each problem may be encoded with a SNOMED code or an ICD code to facilitate recordkeeping and diagnosis, and orders or other procedures associated with each problem also may be encoded with a CPT code to facilitate billing and reimbursement.

[0029] In addition to problem lists, the EHR also may maintain a list of test results for each patient, as well as a history of all procedures recommended to and/or performed on each patient. In the former case, each test result record may be encoded with an ontological code such as a LOINC code. In the latter case, each procedure may be encoded with an ontological code, again, such as a CPT code.

[0030] In each of the instances set forth above, the data in the EHR may be encoded with one or more interface terminology concepts – directly or indirectly. In the case of indirect encoding, the EHR entry may be encoded with an interface terminology description. Rather than being related hierarchically to concepts, each description instead may be an alternative way to express a concept. Each interface terminology concept may be mapped to one or more of the various administrative, clinical, or reference codes, permitting the EHR to retain all of the necessary mapping without devoting memory, storage, or other system resources to maintaining those maps for every single patient record.

[0031] Additionally or alternatively, each computer may be configured to execute patient problem list software and/or access patient problem list data independent of executing EHR software. Such data may reside, for example in a data repository such as a data warehouse (including the problem list software's log repository) or a health information exchange. In this latter alternative, the patient problem list data may originate from one or more sources, including being generated by EHR software, being aggregated from scanned documents via a transcription service, or being received as free-text data input.

[0032] One example of a system and method for maintaining a problem list by mapping to elements of an interface terminology is disclosed in the commonly-owned U.S. patent publication 2015/0242571, published August 27, 2015, the contents of which are incorporated by reference in their entirety.

[0033] Whether the data for the present method originates from patient-centric EHR data or from independent or separate problem list data, that data may include the practitioners' requests and generated response data, e.g., a problem list entry, a procedure entry, etc. That data also may be tagged with enough information to be able to distinguish patients and providers from one another, e.g., each provider may be assigned a unique identifier and each patient may be provided with a separate unique identifier, either on a per-provider or an overall basis. The data will also include a set of problems and the interface terminology problems to which those problems have been associated. Additionally, the data may include demographic information such as age, gender, and location. Still further, the data may include ancillary information such as medications, lab results, etc., in order to provide for more robust analysis among patients. Moreover, additional data such as malformed requests, test requests and improperly entered identifiers (organization, provider, and/or patient) may remain in the data repository but may be ignored.

[0034] Once patient data generally, and patient problem list data in one aspect, specifically, have been encoded with an ontology common to multiple provider systems, it may be possible to identify all problems on a single patient's problem list. Using that same information, the system then may be configured to identify providers who have similar case mixes among their patient sets, while also maintaining confidentiality and privacy of each provider's patient data. As explained in greater detail below, given the input of a plurality of patient problem lists, the system may return a list or graphical display of resulting entity matches, ranked by similarity.

[0035] Additionally, the system may be configured to provide results in multiple use cases. In a first instance, the system may be provider-centric, such that a provider may seek one or more other providers that are similarly matched to the provider's practice, i.e., one or more providers caring for a similar population. This application may be particularly useful to a practitioner diagnosing and/or treating a patient with less commonly-seen symptoms, such that the practitioner may be able to consult with or draw on the experience of one or more relevant providers.

[0036] This application also may be useful to a provider in order to analyze the practitioner's case mix as compared to that of other providers, in order to evaluate one or more patient statistics and to see how those statistics match up against those from the other providers. In one

aspect, a plurality of problem list entries may be used to generate a Hierarchical Condition Category (HCC) or a Diagnosis-Related Group (DRG) code for the patient, and the identity of those codes may vary based on a multiplicity of factors including, e.g., the problems identified as part of the group, the identity of a "primary" versus one or more "secondary" problems, the existence of complications, the presence of co-morbidities, etc. Practitioners, thus, may use the present system to evaluate or compare patients with complex problem lists to determine whether those patients are being documented similarly. Similarly, practitioners may use the present system to determine whether they are making similar referrals to similarly-situated patients.

[0037] In a second instance, the system may be patient-centric, in which a provider may seek out one or more patients within that provider's patient list that have the same or similar problem list as a reference patient. The same HCC and DRG functionality discussed above also may apply in this instance, with the practitioner being able to analyze his or her own patient set for similarities.

[0038] Briefly, and with reference to FIG. 1, the method 100 may include the step of identifying 102 problems on one or more providers' one or more patient lists. In one aspect, the identifying step may include encoding 104 the patient data with a common ontology. From there, the method may extract 106 user or provider identifiers and identifiers for each of the one or more patients. The method then may include the step of parsing 108 the ontology concepts into categories and subcategories. Once those concepts are parsed, the method may include building 110 a comparison model by finding similarly-situated providers or patients, and multiple examples for deriving those similarities are discussed in greater detail below.

[0039] Categorization and Clustering of Problem List Elements

[0040] As seen in FIG. 2, a method for processing electronic medical record problem lists may be employed to generate a clinically relevant patient profile. In one aspect, the patient profile may be useful to a clinician because it may categorize and group related problems according to concept groupings, and groupings may be determined based on semantic distance between the represented concepts. For example, all cardiovascular problems may be grouped under a "cardiovascular" category, all kidney-related problems may be grouped under a "renal" category, etc.

[0041] In addition, the system may attach indicator flags to the problems within each category, which may permit later ranking and ranked display of the problems according to attributes, such as severity, timeliness, or other concepts such as classification within a clinical measure. One example of such a flag is seen in FIG. 2, in which the problem "Diabetes mellitus" and the related problems clustered underneath that summary problem are marked with a CQM flag. The

system may apply an indicator flag to the summary problem if any of its clustered problems (as that term is discussed in greater detail below) include the flag.

[0042] The CQM, i.e., Clinical Quality Measurement, flag indicates that its associated problem element must comply with CQM requirements for treatment and documentation in order to be eligible for the reimbursements provided for such compliance. Thus, a problem having this flag may be presented to the user as a higher value or higher priority problem element. In addition to having the flag callout, this flag also may be used as a factor in problem list ranking. For example, CQM problems may be ranked and presented higher on the problem list within each category than other, non-flagged problem elements.

[0043] Other potential flags may include HCC (Hierarchical Condition Category), CC (Complication and Comorbidity), and MCC (Major Complication and Comorbidity). One of ordinary skill in the art would appreciate that values associated with these terms are reflective of the severity of their underlying problems. As such, problems flagged with one or more of these flags may provide a visual indicator to the user that they may need to be addressed with higher priority than other problems on the list.

[0044] Returning to FIG. 2, multiple criteria in addition to the indicator flags may be applied to the problems in order to determine the rankings within these lists. For example, problems that are associated with/require medication may be ranked higher than those that are/do not. Problems that are entered by a physician/clinician may be ranked higher than those that are sourced from other entities later in the record review process, e.g., by a coder or other administrative personnel. Problems that are obtained from workflow or some other outsider source, e.g., those problems that may be extracted from review of the patient's chart may rank somewhere in between clinician- and coder-generated problems (assuming all other factors are the same). Problem entries may be time-stamped, such that more recent problems may be ranked higher than older problems.

[0045] The system also may generate lists in order to call attention to problems that may require more immediate attention or problems that may affect multiple disciplines. For example, another possible category may be an "in focus now" category, which may display those problems currently most relevant to the user, regardless of whether the problem also can fit into one of the other categories described above, and a "special display" category, which may list high priority problems of extreme, immediate importance, or of problems which are always part of the patient's overall baseline health state. These problems may be categorized more specifically, but they may have effects that cross disciplines, such that the clinician may desire to know about them when addressing the specific problems within his or her discipline.

[0046] This mapping may serve as the basis for the categorization, grouping, rolling up, nesting, etc., of the entries in a problem list. Certain interface terminology concepts may be related to other interface terminology concepts based on similar subject matter. For example, there may be a plurality of concepts that pertain to cardiac conditions. Thus, all problems that map to these concepts may be grouped together for categorization and display such as that shown in FIG. 2.

[0047] In addition to the ranking or sorting criteria describe above, these outside vocabulary mappings may be an additional factor used to rank the problem list entries. For example, mappings to some established terminologies or vocabularies may be used to perform the mapping/grouping described in the previous paragraph, and mappings to a second terminology or vocabulary or a proprietary mechanism may be used to sort more specifically within the determined categories.

[0048] Turning now to FIG. 3, it will be seen that certain problems not only fall within the same category as other problems but that they also may be considered subsets of another problem, i.e., they may be clusters within that problem. These relationships can be determined and managed by using the interface terminology, which also may recognize that certain concepts are more general than others and thus are hierarchically related to those other concepts. The system may group these more specific concepts underneath the more general, parent concept, thereby further arranging the problem list, whose entries may be mapped to these sub-concepts. As it relates to presentation of these problem list entries, the system may display in the problem list the problem that maps to the more general, parent concept and an indicator that other problem entries are nested or clustered and may be viewable under that parent problem, e.g., by clicking on the indicator.

[0049] In one aspect, clustered problem elements underneath a more general, parent concept may be ranked or organized using one or more of the criteria discussed above for ranking elements within the problem list generally. Alternatively, as seen in FIG. 3, clustered problem elements may be arranged using a more simplistic algorithm, e.g., they may be arranged alphabetically. In still another aspect, the system may rank flagged problems above non-ranked problems and then apply the more simplistic algorithm within each of those subsets. In any event, the system may allow user customization, permitting the user to rearrange the ordering of elements both in the problem list and within the clustered subsets, as discussed below.

[0050] From a database management perspective, clustered problems may be stored as a list of elements in a flat file database, with each element pointing to its parent problem element. Alternatively, clusters may be sub-trees in a hierarchical database structure underneath their respective category elements.

[0051] To this point, the patient list has been described as being patient specific, i.e., each patient has his or her own list, with entries specific to that patient in order to accurately record the patient's problem history. The system and method may function similarly as a way to bring a clearer clinical picture for a population aggregator, i.e., determining what problems exist for a given population, or for a given patient who may have multiple problems culled from multiple sources within a large data warehouse. In that case, the number of problems in the aggregated list may be larger (likely significantly larger) than for an individual record within an EHR, although the methodology may remain the same, i.e., each problem may be mapped to an interface terminology concept, concepts may be grouped and ordered, and the ordered problems then may be available for logical display and analysis.

[0052] As seen in FIG. 4, and as discussed above, another issue with problem lists may become evident when attempts are made to combine lists from multiple different sources. These sources may format, store, and/or represent elements in the list differently from one another and not in a consistent format.

[0053] In order to accomplish reconciliation of elements within a single list (i.e., grouping problems within a list into categories and establishing clusters within those categories, which may or may not include the step of combining elements from multiple problem lists into a single list), the system may create an anchoring term from an interface terminology foundation technology that permits creation of a semantic distance between any two other terms from external vocabularies. This anchoring term may be considered a central concept within an interface terminology. In one aspect, determining this anchoring term may be achieved by a concept tagging method, and examples of such a method may be found in the commonly-owned U.S. patent 9,418,150, issued August 16, 2016, the contents of which also are incorporated by reference herein in their entirety.

[0054] Data Structure – Providers

[0055] With reference to FIG. 5, in a provider-centric case, the building step may include, for each provider, determining 112 a number of unique providers, unique clusters, and/or unique categories. The building step then may include tabulating 114 a number of distinct patients having at least one problem in a given set of categories, which may include generating a table containing patient counts, partitioned by cluster (or category) for each provider, and counting/summing the total number of patients for each provider. The building step further may include tabulating 116 a total number of problems for the provider in those categories and calculating 118 a distance in pair-wise comparison from each provider to a plurality of other

providers, and, preferably, to every other provider in order to determine how far away the provider's case mix is from other providers.

[0056] For example, consider a single arbitrary provider, Pr_x , with 4 clusters, i.e., cl_x ; $x = 1; 2; 3; 4$, and 3 patients, i.e., Pa_y ; $y = 1; 2; 3$. The following table depicts the determining step 112, determining how many problems in each cluster a given patient may have in this example:

Patient/Cluster	cl_1	cl_2	cl_3	cl_4
Pa_1	20	32	0	15
Pa_2	19	0	51	73
Pa_3	9	0	15	3

[0057] The next table then depicts the tabulating step 114, counting how many patients have at least one problem in the given clusters. For example, each cell in the following table is the sum of the number of non-zero rows in the respective column of the previous table:

Provider/Cluster	cl_1	cl_2	cl_3	cl_4
Pr_x	3	1	2	3

[0058] The process may be iterated for each provider and each cluster within the data set, yielding a table of provider-cluster count entries.

[0059] Data Structure – Patients

[0060] With reference to FIG. 6, in a patient-centric case, the building step may include, for each provider, determining 120 a number of unique patients, unique clusters, and/or unique categories. The building step then may include sorting 122 each patient's problem list entries into one of the clusters or categories to generate a table containing cluster or category counts for each of the provider's patients. The building step further may include tabulating 124 a total number of problems for each patient in those categories and calculating 126 a distance in pair-wise comparison from each of the provider's patients to each of the provider's other patients in order to determine how far away the patients' problem list mixes are from the other patients.

[0061] For example, for a set of problem clusters for an arbitrary (high-activity) provider, the determining step 120 may result in 25 patients, with problems in 33 clusters and 17 categories.

[0062] For the same provider/cluster/patient combination as in the first example, if any patient went to the provider for a specific problem, the system may consider it as a "1," yielding:

Patient/Cluster	cl ₁	cl ₂	cl ₃	cl ₄
Pa ₁	1	1	0	1
Pa ₂	1	0	1	1
Pa ₃	1	0	1	1

[0063] The system will perform further analytics on this data set in order to determine relative similarities between providers and/or patients, as discussed in greater detail below.

[0064] Data Structure – Partitioned by Time

[0065] Both the provider- and the patient-centric aspects described above may use patient data going back to when the electronic health records or problem list entries first were used to start recording encounter data. Alternatively, the system may parse the records into temporal segments, e.g., monthly segments, such that the matrices depicted above may become three-dimensional, with the third dimension (depth) being time. In that instance, the sum of the depth rows then may be the matrices depicted above.

[0066] In another instance, the system may be configured to let a user select an earliest time before which any data will be excluded, e.g., the past 5 years or 1 year. Still further, the system may be configured to receive a user input of a specific time period from which data is extracted, e.g., between the previous April 1 and June 30.

[0067] Patient Distance Calculation Methodology

[0068] Once the patient data in either aspect described above has been determined, the system may be configured to calculate a count of patients for each provider with at least one problem in each cluster or category.

[0069] Patient counts may fluctuate quite highly from provider to provider and are not bounded above, as the number may increase as the practitioner sees more patients over time. As such, it may be more consistent to normalize the data, e.g., by reporting proportions of patients with problems that present in each cluster rather than absolute counts. Using proportions means that the patient count is always bounded between 0 and 1, inclusive. Additionally, rather than discrete patient count numbers, a proportional representation utilizes a continuous scale.

[0070] Whether using absolute values or proportions, the system then may calculate distance measurements for each provider based on those values or proportions.

[0071] One technique for identifying a distance between two providers is to use a Minkowski distance of order 2 (2-norm). That distance is calculated as the square root of the sum of the differences between cluster or category proportions, i.e.:

$$[0072] d_{ik} = \sqrt{\sum_{i \neq k, j} (P_{ij} - P_{kj})^2}$$

[0073] where:

[0074] d_{ik} is the distance between provider i and provider k ;

[0075] $P_{ij} = \frac{c_{ij}}{c_i}$, i.e., a proportion of patients for provider i with at least one problem in a cluster or category j ,

[0076] c_{ij} is the number of patients for provider I with at least one problem in cluster (category) j ,

[0077] c_i is the total number of patients for provider i .

[0078] In this manner, patient problem lists that include multiple problems within a cluster are weighted more heavily than clusters for which only a single problem appears, reflecting that the former situation may signify a more detailed or complicated case mix.

[0079] The distance calculation may be iterated for each provider pairing, generating a distance matrix:

Provider	Pr ₁	Pr ₂	...	Pr _i	...	Pr _j	...	Pr _n
Pr ₁	0	d_{12}	...	d_{1i}	...	d_{1j}	...	d_{1n}
Pr ₂	d_{21}	0	...	d_{2i}	...	d_{2j}	...	d_{2n}
.
.
.
Pr _i	d_{i1}	d_{i2}	...	0	...	d_{ij}	...	d_{in}
.
.
.
Pr _n	d_{n1}	d_{n2}	...	d_{ni}	...	d_{nj}	...	0

[0080] In this case, $d_{ij} = d_{ji}$, so that the matrix is symmetric.

	[,1]	[,2]	[,3]	[,4]	[,5]	[,6]	[,7]	[,8]	[,9]	[,10]
[1,]	0.0000	1.8048	0.9743	0.5519	0.8130	1.2877	0.7178	0.8683	1.6767	0.6844
[2,]	1.8048	0.0000	2.1985	1.7023	1.6250	1.9365	1.9597	1.7957	2.6458	1.9277
[3,]	0.9743	2.1985	0.0000	0.9877	1.2970	1.4434	1.2047	1.1995	1.1785	1.1837
[4,]	0.5519	1.7023	0.9877	0.0000	0.6621	1.2574	0.9628	0.7858	1.7726	0.5366
[5,]	0.8130	1.6250	1.2970	0.6621	0.0000	1.5052	1.1094	1.0216	2.0143	0.8138
[6,]	1.2877	1.9365	1.4434	1.2574	1.5052	0.0000	1.4743	1.3534	1.9365	1.4604

[7,]	0.7178	1.9597	1.2047	0.9628	1.1094	1.4743	0.0000	1.0880	1.6008	1.0278
[8,]	0.8683	1.7957	1.1995	0.7858	1.0216	1.3534	1.0880	0.0000	1.7555	0.9698
[9,]	1.6767	2.6458	1.1785	1.7726	2.0143	1.9365	1.6008	1.7555	0.0000	2.0123
[10,]	0.6844	1.9277	1.1837	0.5366	0.8138	1.4604	1.0278	0.9698	2.0123	0.0000

[0081] While the distances above were calculated using a 2-norm, it may be possible to calculate those distances in other ways such as, e.g., by using a 1-norm or an infinity- (or max-)norm. In the former case, the 1-norm is a sum of the absolute value of the differences between the P values, and in the latter case, the infinity- (or max-) norm is the maximum absolute difference between the P values.

[0082] Once distances are calculated, clusters may be arranged to indicate similarity between providers. One example of this arrangement may be seen in FIG. 7. In that figure, each circle represents a provider, with the number in the circle corresponding to a provider ID. The various groups of circles reflect clusters.

[0083] Using the relationships shown in FIG. 7, it may be possible to determine which provider(s) is/are most similar to a given provider. Lines connecting one provider to another signify that the linked providers are most similar to one another. Arrows extending from one provider to another reflect directionality. For example, the arrow from provider "20" to provider "86" in the middle of the figure signifies that provider "86" is the closes match to provider "20" but that provider "20" is not necessarily the closest match to provider "86." Due to the nature of the data modeling, each provider may be sent to, at most, one other provider, although a single provider may be sent to by multiple providers.

[0084] In another aspect, the system may receive a provider selection from a user and provide a different graphical representation of the other providers that the system has determined to be most similar based off the distance calculations and/or the clustering of providers. For example, FIG. 8 depicts one example of this representation, with the selected provider placed at the center of a provider mapping and the related providers displayed as radiating outward from that center. Related providers may be joined to the selected provider, e.g., via a line connecting provider IDs. Similarity may be represented by a thickness of the connecting line, with a thicker line signifying more closely-related providers. Additionally or alternatively, similarity also may be reflected by radial proximity of the other providers to the selected provider in the display.

[0085] In another aspect, the system may be configured to use categories to calculate the distance matrix referred to above. Using these categories, the same general methodology set forth above may be followed, substituting clusters for categories. Since the list of problems in a

category may not overlap completely with the list of problems in a cluster, the underlying problem counts may vary. As such, using categories instead of clusters may result in the same providers being mapped to different other providers, changing the display of logical relationships, as seen in FIG. 9.

[0086] Edges

[0087] The analysis may be refined further by applying a weight to the distance matrices based on cluster and category, generating a modified distance matrix according to the following formula:

$$\mathbf{[0088]} \quad d_{new} = w \times d_{category} + (1 - w) \times d_{cluster},$$

[0089] where $w \in [0,1]$. The selection for the weight, w , may be user-defined in order to account for whether the weight should be more specific or more generic.

[0090] In order to provide more meaningful results to the user, the system may filter the results to a predetermined or user-selected number upon selection of a provider for which greater analysis is desired. For example, upon selection of provider number 1, FIG. 10 is one example of a graphical depiction of the twenty providers most similar to that selected provider, displayed as a similarity group surrounding that provider. It will be appreciated that the grouping may reflect more than just the providers provided as direct elements of a cluster as in FIGS. 4 and 6.

[0091] In still another aspect, the system may be configured to filter providers according to a temporal component of activity. For example, as seen in FIG. 11, only providers with relevant patient activity in the past predetermined or user-selected time period, e.g., the past month, may be used for comparative purposes. Alternatively, the system may accept user input of a different specific time period, e.g., the time span between 30 and 60 days ago. In either case, this filtering capacity may allow the user to search for providers with relevant activity either most recently or within the desired time period. In the former case, the filtering may be useful to filter out providers that have not had any relevant activity recently and, as such, may not be as useful. In the latter case, the filtering may be useful when a historical analysis is desired, such as determining a temporal or geographical origin or spread of particular disease.

[0092] Providers for which no relevant activity was recorded in the given time period may be hidden from presentation to the user. Alternatively, those other providers may be represented in the display, but they may be not connected to any other provider, reflecting their lack of commonality with the other providers.

[0093] Patient Analysis

[0094] In addition to analyzing similarity on a provider basis, the system also may be configured to analyze patient similarity. Each patient may be assigned to one or more clusters, the clusters

represented in a database table with one or more unique identifiers such as as one or more cluster_title entries. Patient similarity analysis using those cluster_title variables may be accomplished in the same way as the provider analysis set forth above, and results may be displayed in a similar fashion as see, e.g., in FIG. 12.

[0095] Other patient analysis, using other patient-centric variables, may be carried out in the same way. Alternatively, as with the provider analysis described above, patient-centric analysis may be done using categories instead of clusters, with the results similarly changing, as depicted in FIG. 13.

[0096] In still another aspect, as seen in FIG. 14, patients may be ranked depending on whether they have the exact same problems, as opposed to just similar ones. As seen in that figure, the similarity scores may be identical for multiple other patients, reflecting the problem match.

[0097] Turning to FIG. 15, it further may be possible to analyze patients or providers based on a combination of problems and clusters. One or more problems and one or more clusters may be evaluated. Additionally, it may be possible to provider user weighting for one or more of the problems or clusters, e.g., if it is deemed that the presence of a particular problem or combination of problems is more significant than the other variables. In one aspect, any weight may be provided to each variable, provided that a sum of the weights is 1. For example, FIG. 15 reflects the use of a single problem and a single cluster variable, each weighted equally.

Conversely, FIG. 16 reflects the use of the same problem and cluster variables, but with the problem weighted at 0.8 and the cluster weighted at 0.2. Still further, FIG. 17 reflects the use of the same problem and cluster variables, but with the problem weighted at 0.2 and the cluster weighted at 0.8.

[0098] In another aspect, any weights can be provided to each variable, and the system will normalize them in order to determine the relative weights as among each variable.

[0099] Additionally, it may be possible to refine the results by reducing the number of clusters being analyzed. For example, a data set may include 120 unique clusters as between 112 unique providers. Of those, 23 clusters may not include any relevant patient data, so that eliminating them from the analysis may yield 97 clusters.

[00100] Still further, the system may be configured to keep or eliminate clusters above or below a user-defined or predetermined correlation. For example, to avoid multicollinearity, or clusters that are highly correlated, the system may eliminate clusters with a maximum correlation above 0.7. In this example, this may reduce the number of clusters from 97 to 61, or essentially halving the number of clusters to be analyzed.

[00101] While it may seem counterintuitive to eliminate highly correlated clusters, the correlation may, in fact, represent an over-weighting of dependent information. For example, problems may be assigned to multiple clusters, e.g., a weight-related cluster and a diabetes cluster, where much of the information or many of the problems in one cluster also apply to the other cluster. In that case, similarity analysis using both clusters may, in effect, double-count those problems, skewing the impact the problems have on the distance analysis. Correlation determinations may be user-defined in one aspect and automatically determined, such as by comparing a number of overlapping problems in each cluster, a proportion of overlapping problems as compared to a total number of problems in each cluster, or via another algorithm, as would be appreciated by one of ordinary skill in the relevant art.

[00102] Outputs may include:

[00103] One or more indicators of similarity among providers based on clusters or categories, overall and/or within a specific set (e.g., a single account).

[00104] One or more indicators of similarity among the providers based on the combination of clusters and categories.

[00105] One or more indicators of similarity as above, partitioned over months or some other predefined or user-defined time segment.

[00106] A ranking of a predetermined or user-defined, e.g., top 20, providers similar to a specific provider based on the combination of clusters and categories.

[00107] A ranking of patients similar to a specific patient within a single provider based on clusters (category).

[00108] A ranking of patients similar to a specific patient within a single provider based on the exact problems, as reflected by encoding of those problems using interface terminology concepts or descriptions, and the clusters.

[00109] The system and method described herein may provide a user with one or more of the following benefits:

[00110] Identifying the single-most similar provider or a plurality of similar providers based on the problem lists associated with one or more of the providers' patients. This identification may be accomplished through a user interface that presents the relationships in a manner similar to FIGS. 7-17, where selection of a provider or patient may cause the display to be refreshed or updated with more specific information about that provider or patient, and/or with additional information explaining the basis for the similarity with that provider or patient and its upstream and/or downstream matches. The identification alternatively may be accomplished by providing a user interface that includes a list or table of either all providers or

patients or, instead, of a predetermined or user-selected number of the closest matches. The list or table may be sortable, e.g., by closest match, and also may be selectable, such that additional information such as the basis for the similarity matches may be displayed upon user selection of a provider or patient. For patient matches, the user interface also may be configured such that, upon selection of a patient identifier, both the selected patient and the subject patient's problem lists may be displayed to facilitate comparison of the two.

[00111] Given a provider, the most similar providers can be determined and provided to the user as an ordered list, e.g., within an organization or a set of organizations, or over all users or providers for which the system has the necessary data.

[00112] An ability to rank providers based on counts or proportions of patients, with problems partitioned by one or more of clusters and categories.

[00113] Within a provider's data set, the same analysis can be performed on a patient-centric basis.

[00114] While the foregoing written description of the invention enables one of ordinary skill to make and use what is considered presently to be the best mode thereof, those of ordinary skill will understand and appreciate the existence of variations, combinations, and equivalents of the specific exemplary embodiment and method herein. The invention should therefore not be limited by the above described embodiment and method, but by all embodiments and methods within the scope and spirit of the invention as claimed.

We claim:

1. A method for extracting data from a plurality of electronic data repositories to provide provider and patient data similarity scoring, comprising:
 - encoding a problem list for a plurality of patients with concepts from a common electronic health record ontology, wherein the plurality of electronic data repositories are maintained by a plurality of providers;
 - parsing the concepts into a plurality of clusters or categories;
 - determining, for each of the providers, a total number of patients that have at least one problem in a cluster or category;
 - iterating the determining step for each of the remaining clusters or categories; and
 - calculating, for each pair of providers, a distance between the providers.
2. The method of claim 1, wherein the calculating step includes calculating a Minkowski distance of order 2 between each pair of providers.
3. The method of claim 1, wherein the calculating step includes normalizing the results of the determining and iterating steps by dividing those results by a respective provider's total number of patients.
4. The method of claim 1, wherein the common electronic health record ontology is an interface terminology.
5. The method of claim 4, wherein the problem list also is encoded, directly or indirectly, with a health record ontology selected from the group consisting of the Systematized Nomenclature of Medicine and the International Classification of Disease.
6. The method of claim 1, wherein the calculating step uses both the results of the determining and iterating steps as well as exact problem list matches as inputs.
7. The method of claim 6, wherein the results of the determining and iterating steps and the exact problem list matches are weighted equally.

8. The method of claim 6, wherein the results of the determining and iterating steps are weighted more heavily than the exact problem list matches.
9. The method of claim 6, where the exact problem list matches are weighted more heavily than the results of the determining and iterating steps.
10. The method of claim 1, wherein at least one of the electronic data repositories is an electronic health record repository.
11. The method of claim 1, wherein at least one of the electronic data repositories is a problem list log repository.
12. A method for extracting data from an electronic data repository to provide provider and patient data similarity scoring, comprising:
 - encoding a problem list for a plurality of patients with concepts from a common electronic health record ontology;
 - parsing the concepts into a plurality of clusters or categories;
 - determining, for each patient, which of the plurality of clusters or categories correspond to at least one concept encoded in the patient's problem list; and
 - calculating, for each pair of patients, a distance between the patients.
13. The method of claim 12, wherein the calculating step includes calculating a Minkowski distance of order 2 between each pair of patients.
14. The method of claim 12, wherein the calculating step includes normalizing the results of the determining step by dividing those results by a respective patient's total number of problems.
15. The method of claim 12, wherein the common electronic health record ontology is an interface terminology.
16. The method of claim 15, wherein the problem list also is encoded, directly or indirectly, with a health record ontology selected from the group consisting of the Systematized Nomenclature of Medicine and the International Classification of Disease.

17. The method of claim 12, wherein the calculating step uses both the results of the determining step as well as exact problem list matches as inputs.
18. The method of claim 17, wherein the results of the determining step and the exact problem list matches are weighted equally.
19. The method of claim 17, wherein the results of the determining step are weighted more heavily than the exact problem list matches.
20. The method of claim 17, where the exact problem list matches are weighted more heavily than the results of the determining step.
21. The method of claim 12, wherein the electronic data repository is an electronic health record repository.
22. The method of claim 12, wherein the electronic data repository is a problem list log repository.

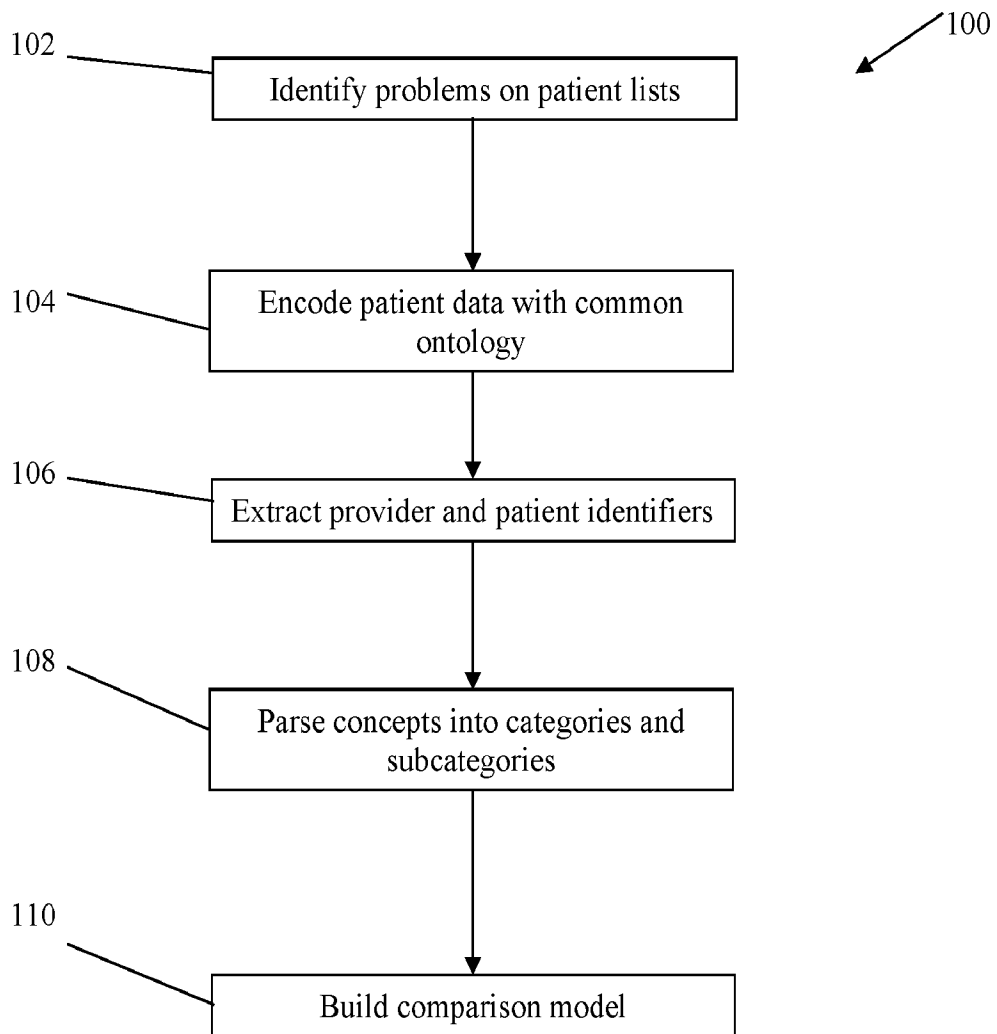

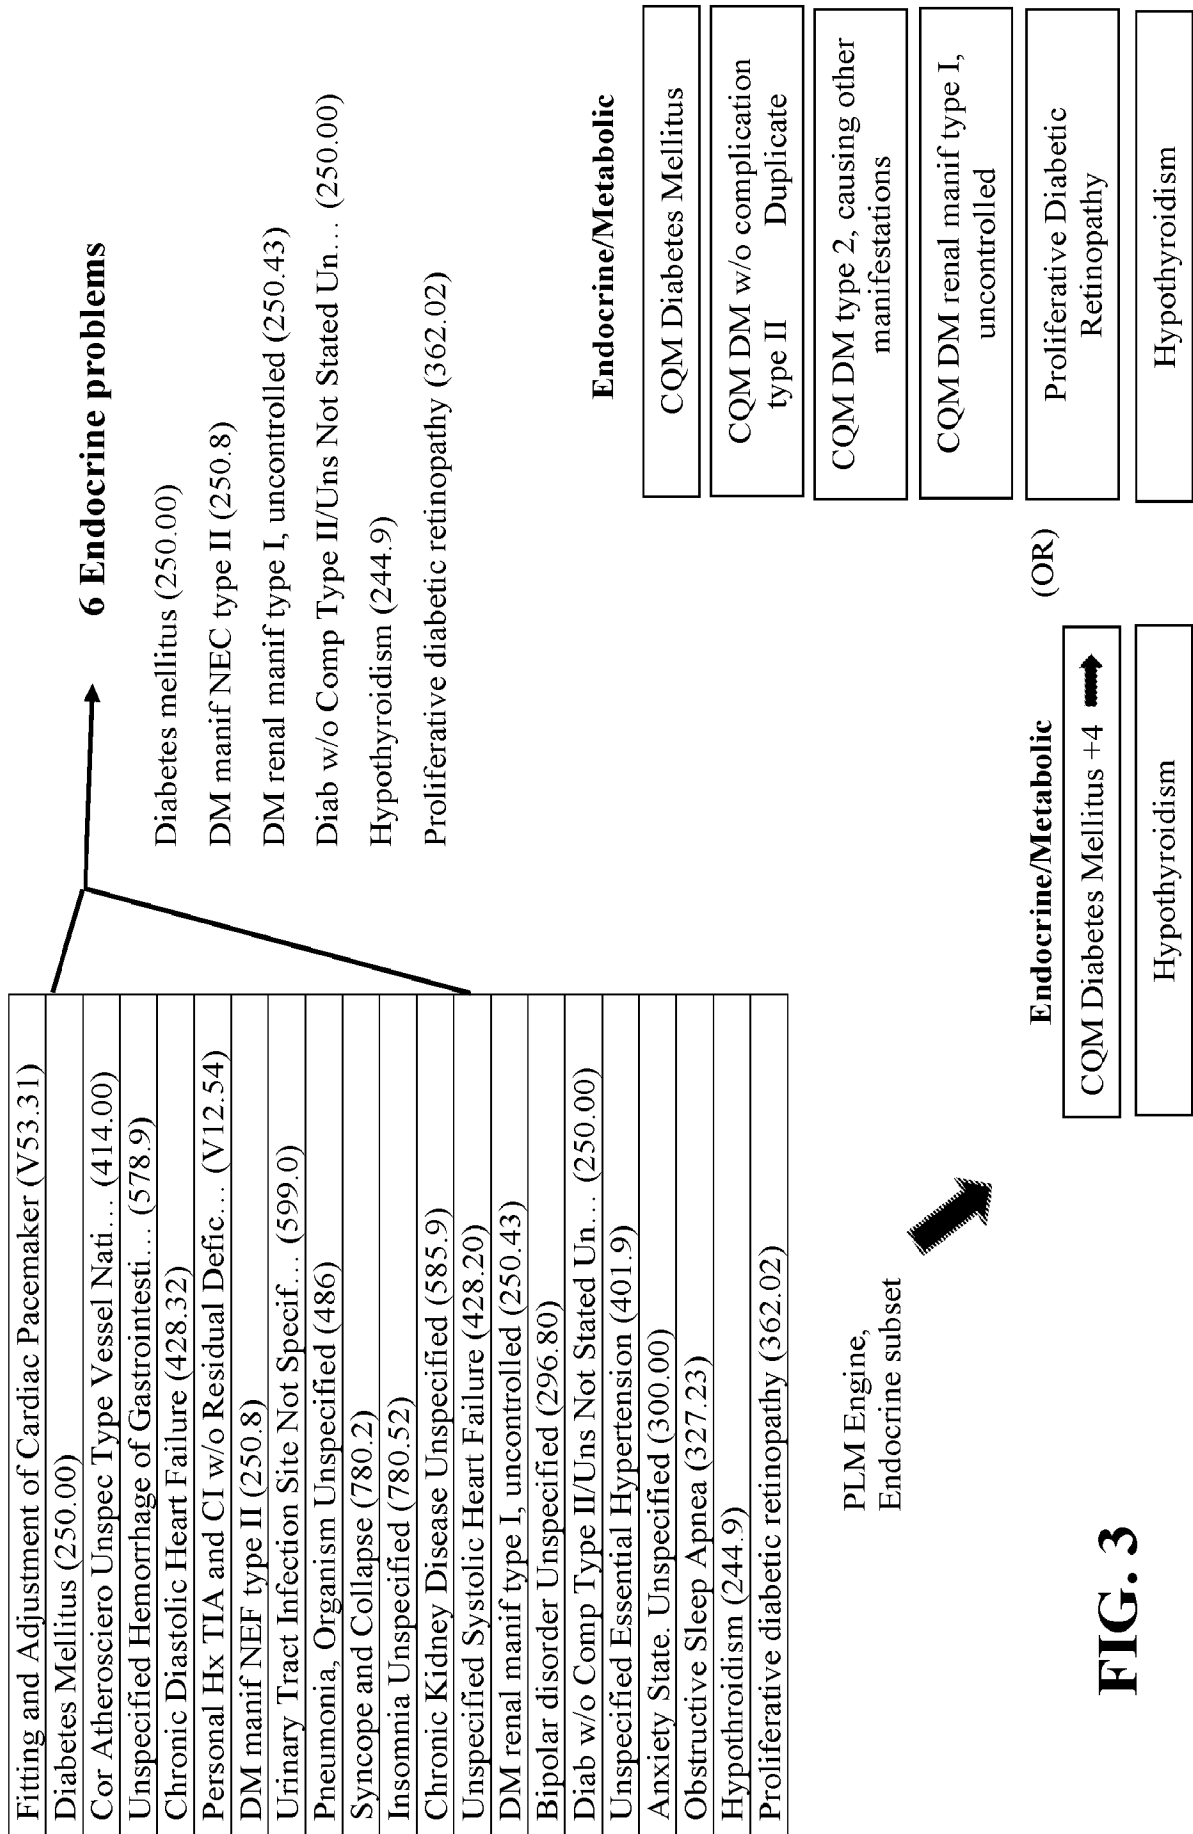


FIG. 1

FIG. 2

PROBLEM 1		PROBLEM 21
PROBLEM 2		PROBLEM 22
PROBLEM 3		PROBLEM 23
PROBLEM 4		PROBLEM 24
PROBLEM 5		PROBLEM 25
PROBLEM 6		PROBLEM 26
PROBLEM 7		PROBLEM 27
PROBLEM 8		PROBLEM 28
PROBLEM 9		PROBLEM 29
PROBLEM 10		PROBLEM 30
PROBLEM 11		PROBLEM 31
PROBLEM 12		PROBLEM 32
PROBLEM 13		PROBLEM 33
PROBLEM 14		PROBLEM 34
PROBLEM 15		PROBLEM 35
PROBLEM 16		PROBLEM 36
PROBLEM 17		PROBLEM 37
PROBLEM 18		PROBLEM 38
PROBLEM 19		PROBLEM 39
PROBLEM 20		PROBLEM 40





PROBLEM 1	PROBLEM 21
PROBLEM 2	PROBLEM 22
PROBLEM 3	PROBLEM 23
PROBLEM 4	PROBLEM 24
PROBLEM 5	PROBLEM 25
PROBLEM 6	PROBLEM 26
PROBLEM 7	PROBLEM 27
PROBLEM 8	PROBLEM 28
PROBLEM 9	PROBLEM 29
PROBLEM 10	PROBLEM 30
PROBLEM 11	PROBLEM 31
PROBLEM 12	PROBLEM 32
PROBLEM 13	PROBLEM 33
PROBLEM 14	PROBLEM 34
PROBLEM 15	PROBLEM 35
PROBLEM 16	PROBLEM 36
PROBLEM 17	PROBLEM 37
PROBLEM 18	PROBLEM 38
PROBLEM 19	PROBLEM 39
PROBLEM 20	PROBLEM 40

- ▶ Septic arthritis of knee, right
- ▶ Lymphedema of lower extremity
- ▶ Chronic Kidney Disease, stage III (moderate)
- ▶ Hyperparathyroidism due to renal insufficiency
- ▶ Hypertensive disorder
- ▶ History of kidney transplant
- ▶ Infected traumatic ulcer of lower extremity
- ▶ Anemia in chronic kidney disease
- ▶ Pseudomonas aeruginosa infection




FIG. 4

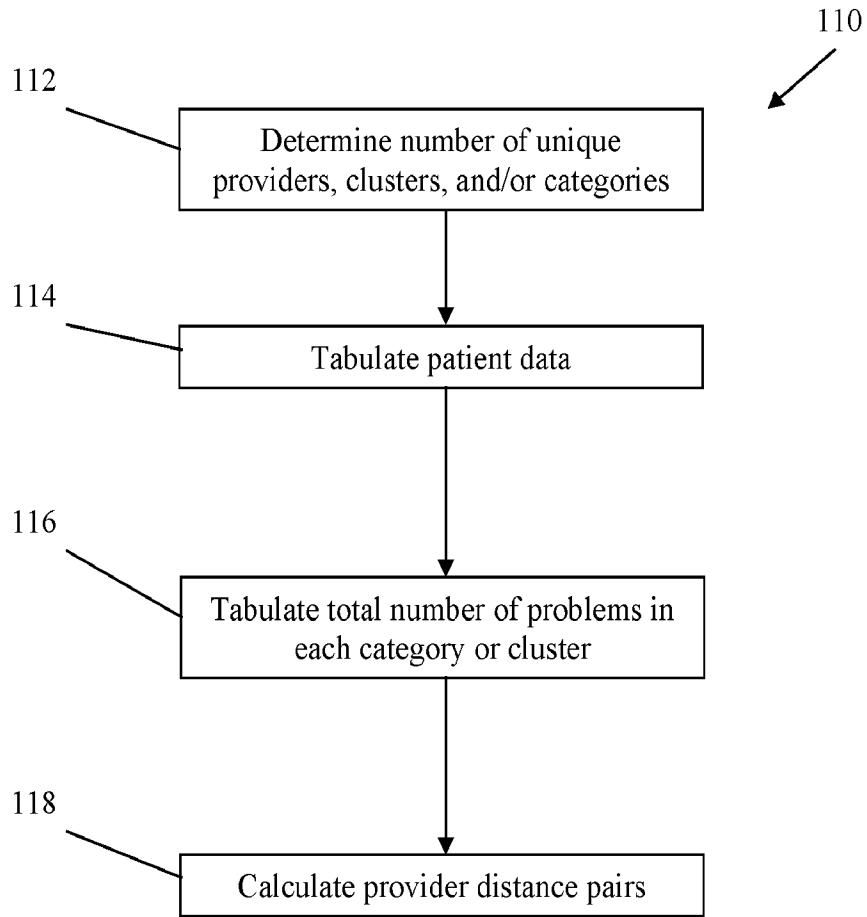


FIG. 5

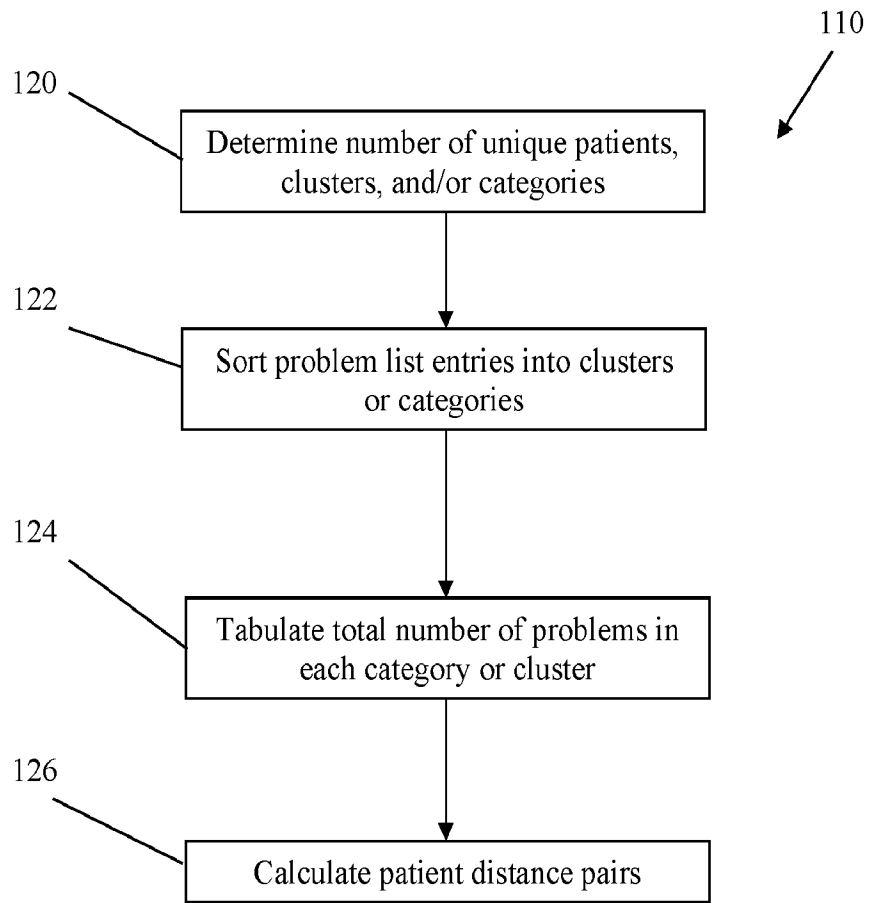


FIG. 6




FIG. 7

> rank top			
	top_prov	Distance	rank
[1,]	10	0.0000000	1.0
[2,]	14	1.0900787	18.0
[3,]	16	0.8461970	8.5
[4,]	22	0.9686442	16.5
[5,]	24	1.1276546	20.5
[6,]	25	1.1276546	20.5
[7,]	26	0.9686442	16.5
[8,]	41	0.8461970	8.5
[9,]	47	0.3918538	3.0
[10,]	53	0.9094836	13.5
[11,]	59	0.8461970	8.5
[12,]	64	0.8461970	8.5
[13,]	69	0.8461970	8.5
[14,]	71	0.4070653	4.0
[15,]	72	1.1276546	20.5
[16,]	73	0.8461970	8.5
[17,]	83	0.8148411	3.5
[18,]	87	0.9094836	13.5
[19,]	97	0.9094836	13.5
[20,]	100	0.9094836	13.5
[21,]	109	0.3422872	2.0
[22,]	112	1.1276546	20.5

Provider no. 10

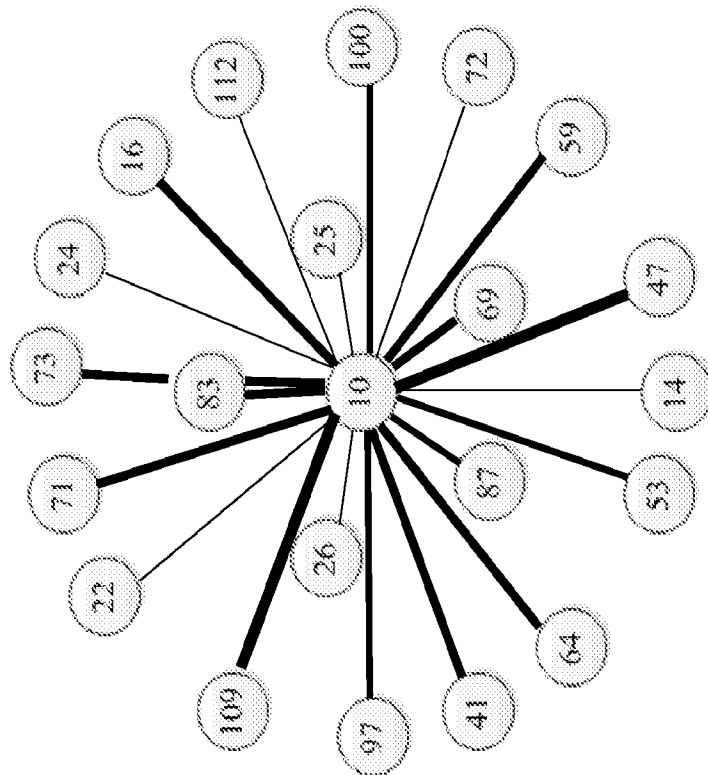


FIG. 8

Using proportion for category_top




FIG. 9

graph of Category_proportion for provider1_top20




FIG. 10

Using proportion for category_top_sub_April





FIG. 11


graph for patient_1 using cluster



[1]	0.000000	3.000000	3.000000	4.358899	4.123106	3.316625
[7]	3.316625	4.690416	3.316525	4.000000	3.464102	5.000000
[13]	3.000000	2.645751	2.828427	3.162278	4.000000	4.472136
[19]	5.099020	6.928203	3.316625	3.000000	3.000000	3.000000
[25]	3.316625					

FIG. 12


graph for patient_1 using category



[1]	0.000000	3.316625	3.316625	4.795832	4.358899	3.872983
[7]	3.872983	6.782330	3.872983	4.472136	3.464102	5.744563
[13]	3.605551	2.645751	2.828427	3.464102	3.741657	4.898979
[19]	5.477226	7.348469	3.872983	3.316625	3.316625	3.316625
[25]	3.872983					

FIG. 13


graph of lexical for patient_1



[1]	0.000000	2.828427	2.828427	2.828427	2.449490	2.828427
[7]	2.828427	2.828427	2.828427	2.828427	2.828427	2.645751
[13]	2.828427	2.645751	2.449490	2.828427	2.645751	2.828427
[19]	2.828427	2.645751	2.828427	2.828427	2.645751	2.828427
[25]	2.828427					

FIG. 14


graph of 0.5*lexical+0.5*cluster for patient_1



[1]	0.000000	5.828427	5.828427	7.187326	6.572595	6.145052
[7]	6.145052	7.518843	6.145052	6.828427	6.292529	7.645751
[13]	5.828427	5.291503	5.277917	5.990705	6.645751	7.300563
[19]	7.927447	9.573955	6.145052	5.828427	5.645751	5.828427
[25]	6.145052					

FIG. 15


graph of $0.8 \cdot \text{lexical} + 0.2 \cdot \text{cluster}$ for patient_1



[1]	0.000000	2.862742	2.862742	3.134521	2.784213	2.926067
[7]	2.926067	3.200825	2.926067	3.062742	2.955562	3.116601
[13]	2.862742	2.645751	2.525277	2.895197	2.916601	3.157169
[19]	3.282546	3.502242	2.926067	2.862742	2.716601	2.862742
[25]	2.926067					

FIG. 16

graph of $0.2 \cdot \text{lexical} + 0.8 \cdot \text{cluster}$ for patient_1



[1]	0.000000	2.965685	2.965685	4.052805	3.788382	3.218985
[7]	3.218958	4.318018	3.218985	3.765685	3.336967	4.529150
[13]	2.965685	2.645751	2.752640	3.095508	3.729150	4.143394
[19]	4.644901	6.071713	3.218985	2.965685	2.929150	2.965685
[25]	3.218985					

FIG. 17

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 17/59025

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06Q 50/00 (2017.01)

CPC - G06Q 50/24, G06Q 50/22, G06F 19/322, G06Q 40/08, G06Q 10/10, G06Q 30/02, G06Q 10/10
G06Q 10/06, G06Q 30/06, G06Q 30/0601

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

See Search History Document

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

See Search History Document

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

See Search History Document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2015/0356270 A1 (Devarakonda et al.) 10 December 2015 (10.12.2015), entire document, especially paras [0014]-[0022], [0029]-[0032], [0074]	1-22
Y	US 2013/0268532 A1 (Doshi) 10 October 2013 (10.10.2013), entire document, especially paras [0013], [0069]	1-22
A	US 2011/0288877 A1 (Ofek et al.) 24 November 2011 (24.11.2011), entire document.	1-22

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

07 January 2018 (07.01.2018)

Date of mailing of the international search report

25 JAN 2018

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450
Facsimile No. 571-273-8300

Authorized officer:

Lee W. Young

PCT Helpdesk: 571-272-4300
PCT OSP: 571-272-7774