



(19) **United States**

(12) **Patent Application Publication**
PATRICK et al.

(10) **Pub. No.: US 2018/0189381 A1**

(43) **Pub. Date: Jul. 5, 2018**

(54) **FRAMEWORKS AND METHODOLOGIES FOR ENABLING SEARCHING AND/OR CATEGORISATION OF DIGITISED INFORMATION, INCLUDING CLINICAL REPORT DATA**

Publication Classification

(51) **Int. Cl.**
G06F 17/30 (2006.01)
G16H 10/60 (2006.01)
G06F 17/27 (2006.01)
(52) **U.S. Cl.**
CPC *G06F 17/30634* (2013.01); *G06F 17/2735* (2013.01); *G06F 17/30705* (2013.01); *G16H 10/60* (2018.01)

(71) Applicant: **HEALTH LANGUAGE ANALYTICS PTY LTD, Eveleigh (AU)**

(72) Inventors: **Jon PATRICK, Sydney (AU); Pooyan ASGARI, Sydney (AU); Min LI, Sydney (AU)**

(57) **ABSTRACT**

The present disclosure relates to frameworks and methodologies for enabling categorisation and/or searching of digitised information, including clinical report data. Embodiments of the invention have been particularly developed to assist categorisation of digitised information, such as clinical report data, in a streamlined manner based on a pre existing set of classification codes. This, in some embodiments, enables the discovery and extraction of meaningful patterns from unstructured clinical reports. Further embodiments of the invention have been particularly developed to assist in the discovery and extraction of meaningful patterns from an unstructured set of digitised information such as unstructured clinical reports. While some embodiments will be described herein with particular reference to those applications, it will be appreciated that the invention is not limited to such a field of use, and is applicable in broader contexts.

(21) Appl. No.: **15/740,218**

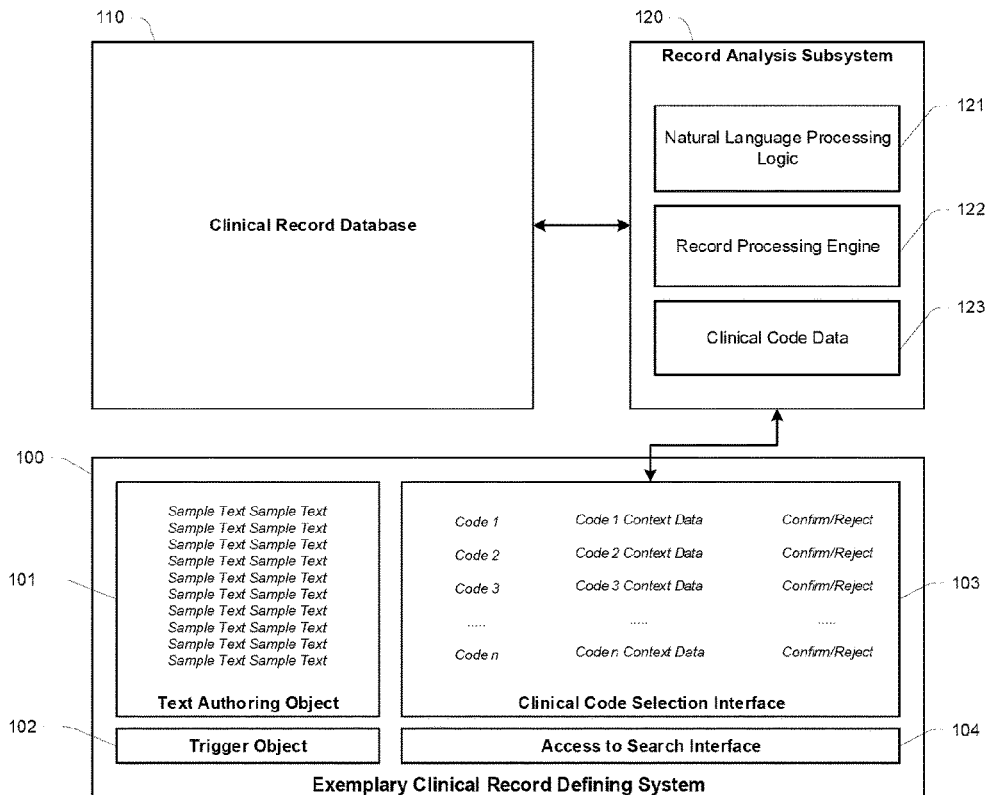
(22) PCT Filed: **Jun. 30, 2016**

(86) PCT No.: **PCT/AU2016/000235**

§ 371 (c)(1),
(2) Date: **Dec. 27, 2017**

(30) **Foreign Application Priority Data**

Jun. 30, 2015 (AU) 2015902540
Jun. 30, 2015 (AU) 2015902542



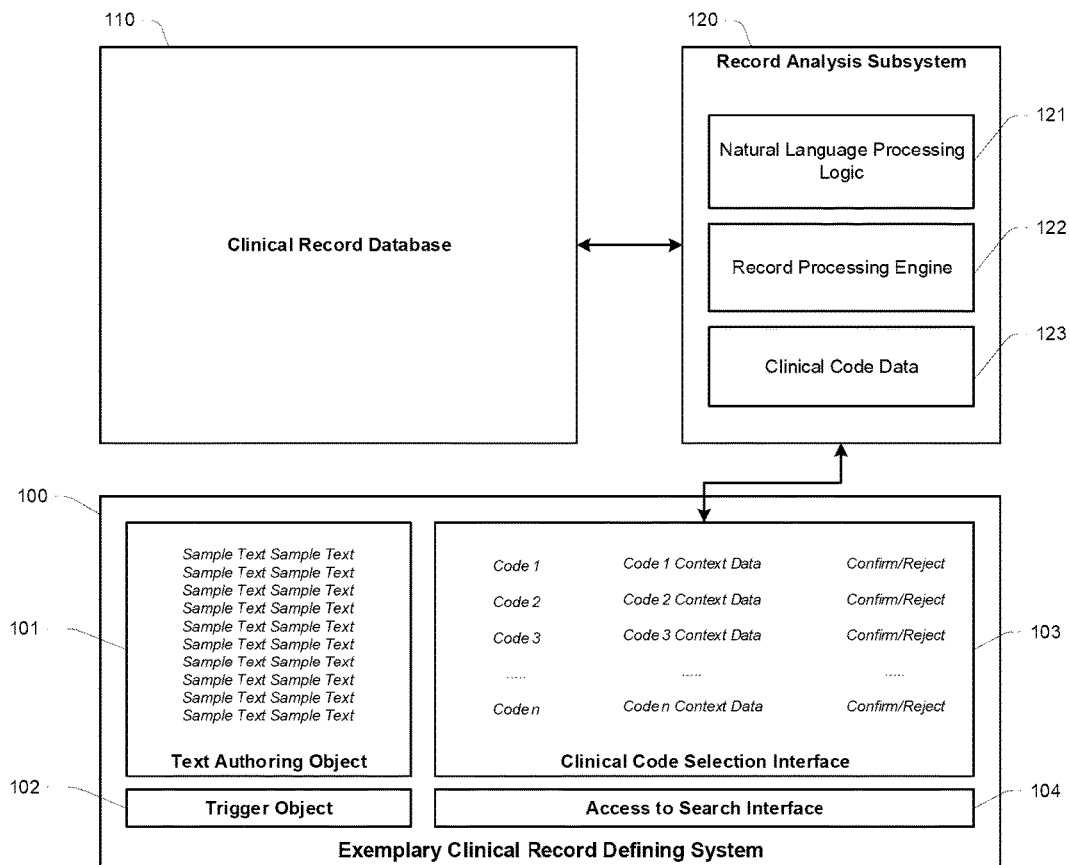


FIG. 1

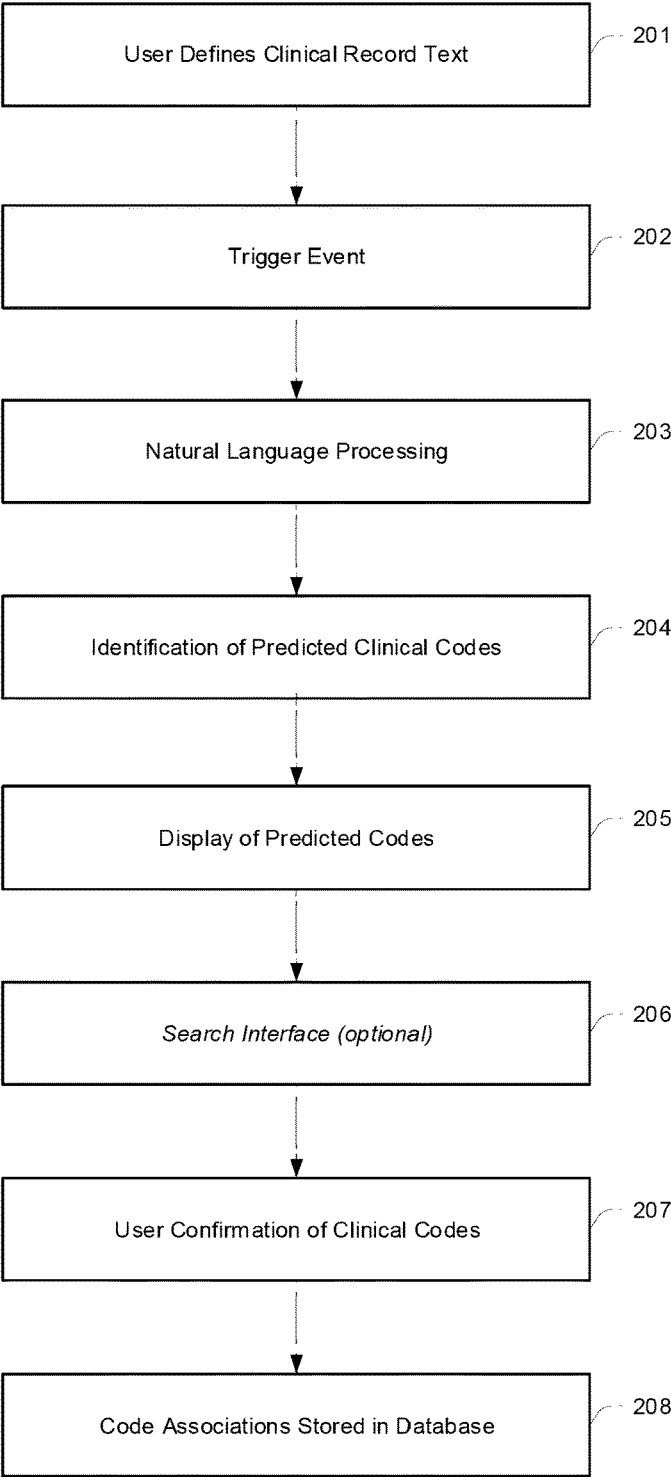


FIG. 2

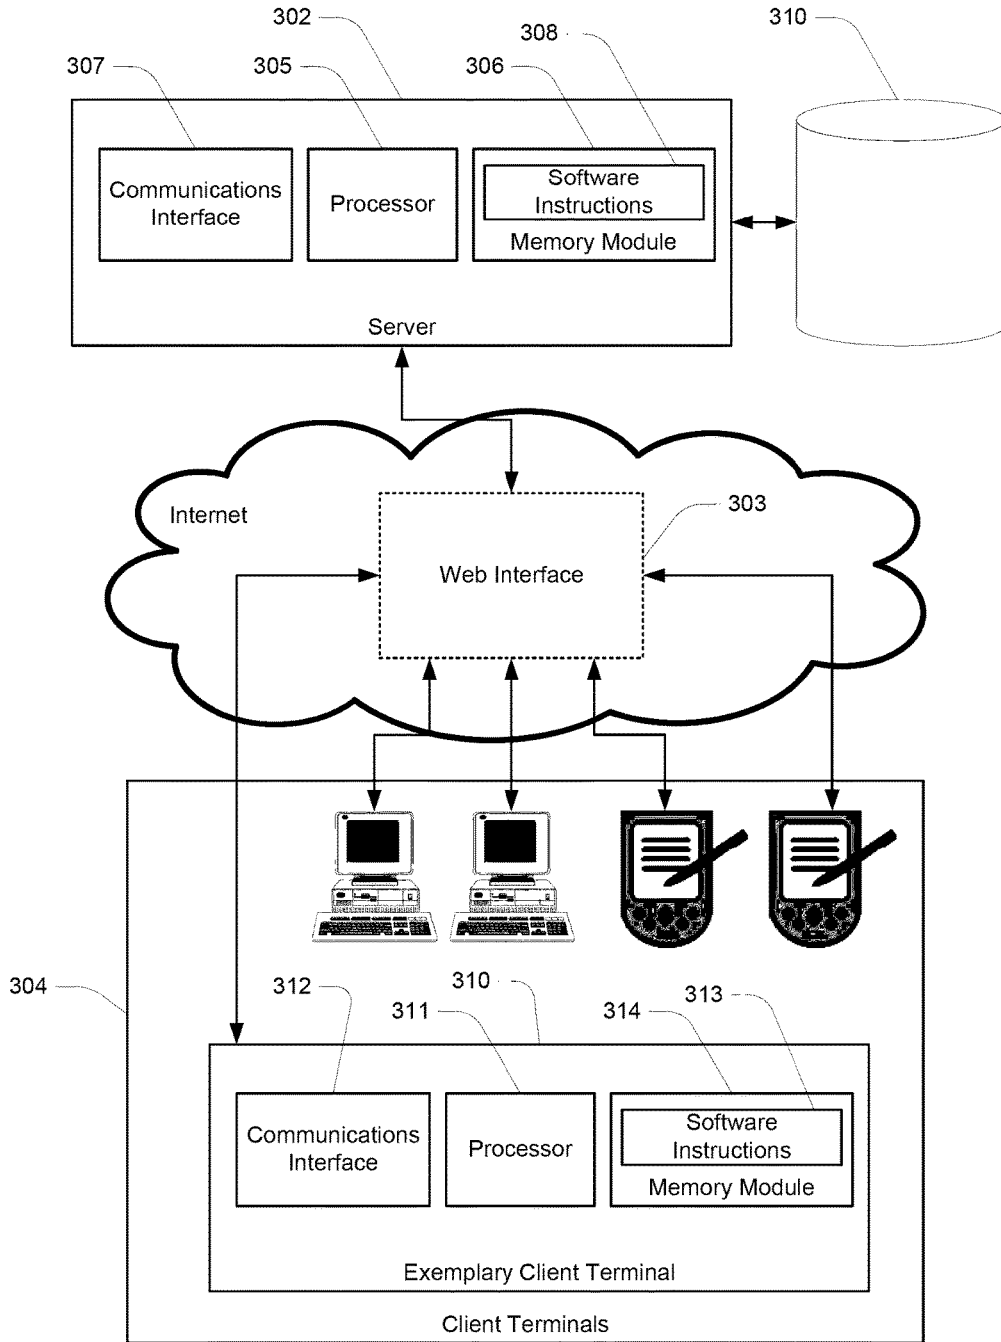


FIG. 3

Click Client

View Details

Text summary	Table of contents	Structure summary	Summary	Text	NCI Description	NCI Code	Context
FP	Section	Disease Class					
1 / 4	F	Other	Respiratory system disease	Non-small cell lung cancer	Non-small cell lung cancer	U449700	Phase I/II BOLN-8001-010, PEG-MAN Non-small cell lung cancer Efficacy combined with report for C
Click to Deactivate	2 / 4	Clinical History	Treatment	Biopsy	Biopsy	8437004	Indication included: Pleural biopsy in biopsy for metastatic
	3 / 4	Clinical History	Respiratory system disease	non-small cell lung cancer	Non-small cell lung cancer	U449700	Pleural biopsy is compatible with non-small cell lung cancer biopsy phase I/II BOLN-8001
	4 / 4	Clinical History	Respiratory system disease	COAD	Colon adenocarcinoma disease	U564100	Indication: 66 year old female COAD Chromosomes: CT, DD, DD, D

Done (1/2/1)

Print Export Selected

FIG. 4

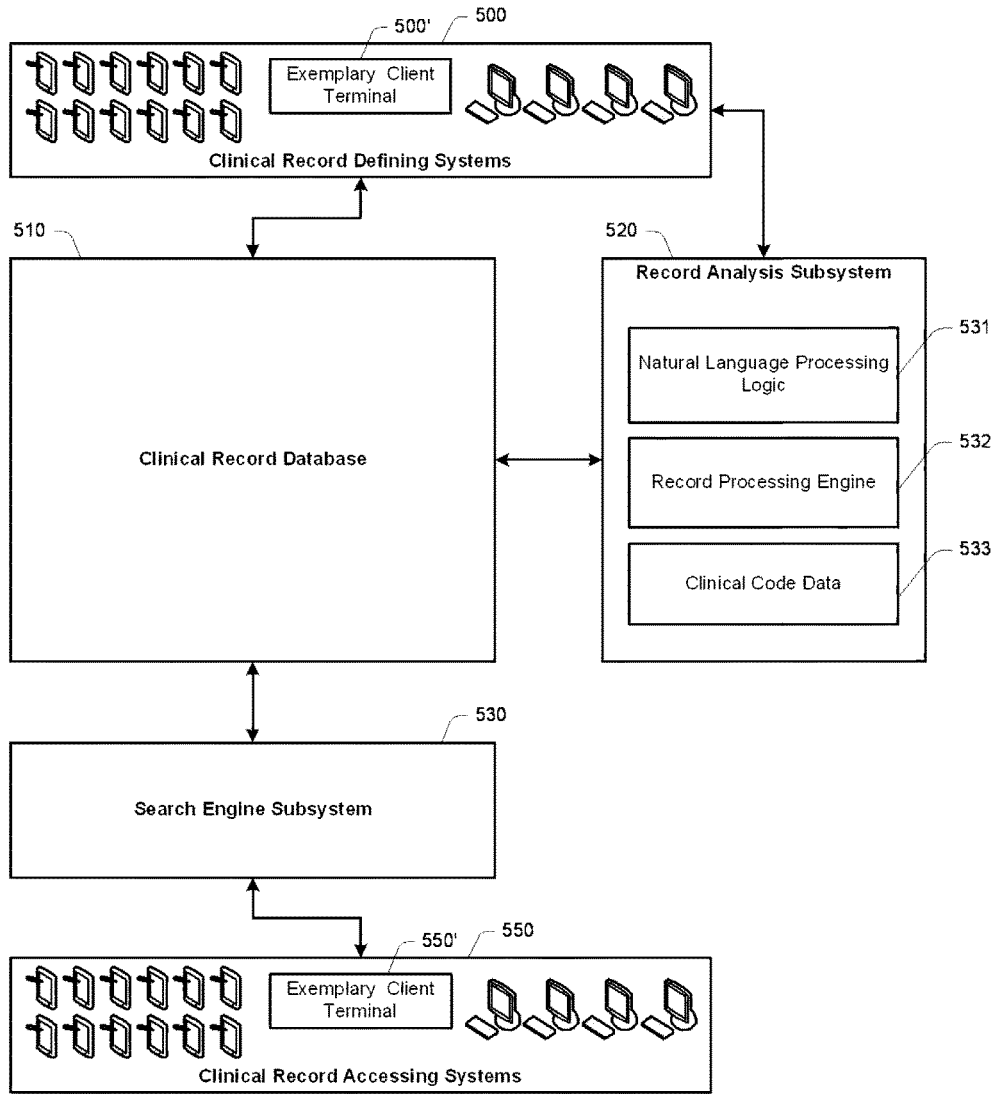


FIG. 5

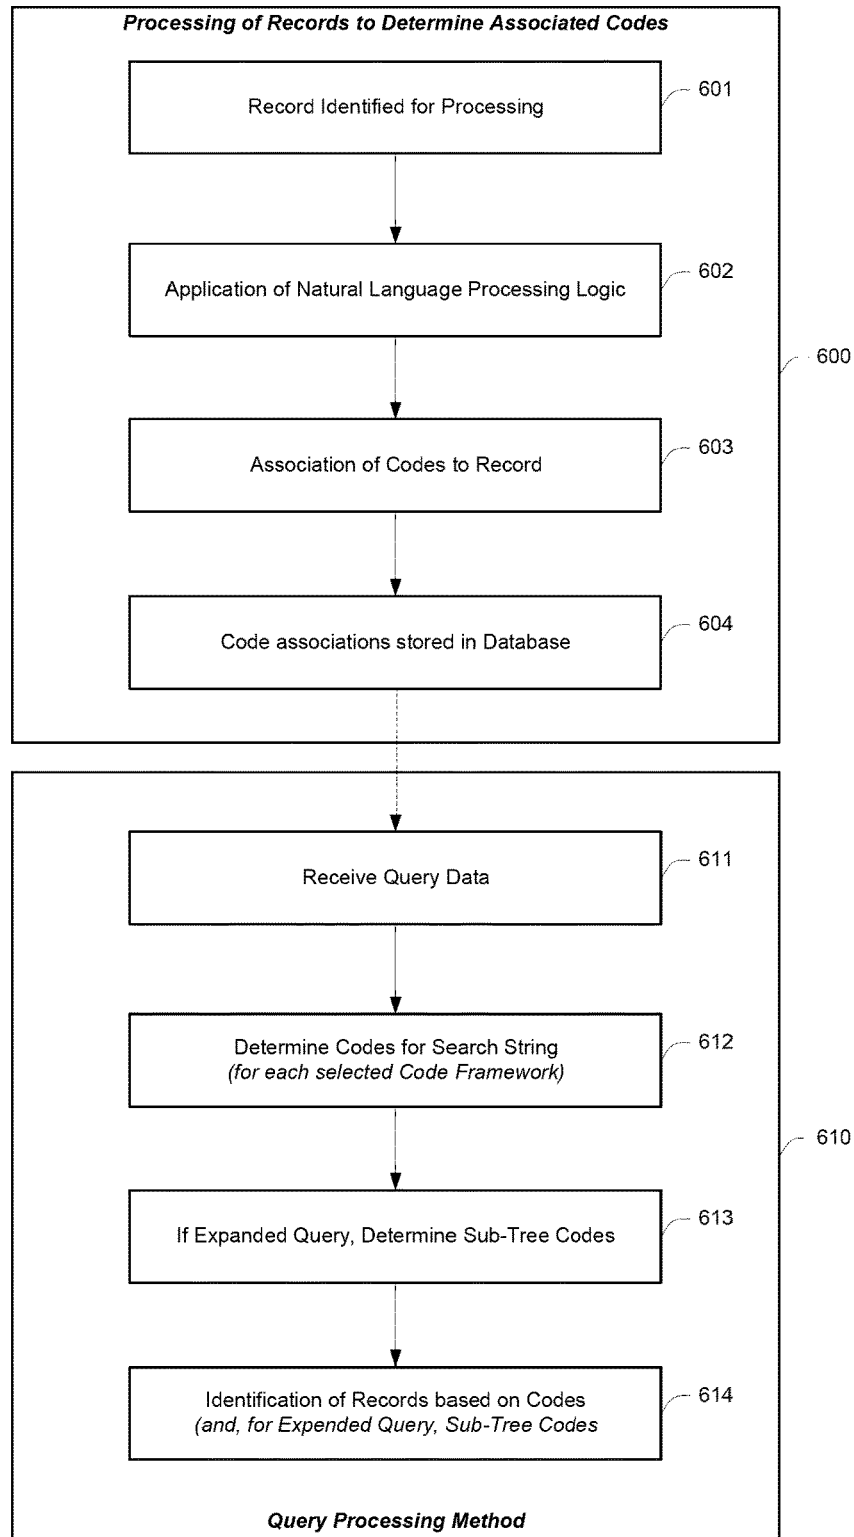


FIG. 6

CliniSearch Results: 34 GO 100 11 Active Database: Radiology_1 FFilename: 326186273.am

Keyword / Phrase: Analytics Engine: Findings:

Terminology Code or Description:

Terminology Settings: ICD-10 SNOMED Search includes subtree

Search Logic Settings: All Any

Refine Method: Constrain Exclude Extend

Refine Target: Current Result Between Saved Results Refine

Search Results: Matched with 34 records

File Name	Hits	Context of Sample Hits from Reports
10 326192599...	1	with Hims with Patient Clinical indications: Invasive ductal carcinoma left breast at 2 & 3 o'clock. 2nd lesion ad, ad, ce
11 326188212...	1	R how/ab/sma Clinical indications: Past history of breast cancer. Back pain. CT - PELVIS Findings: Bilateral sd
12 326192233...	1	dermatomas. Further LC'S and a small focus of invasive lobular carcinoma. Negative nodes on sentinel node biopsy. Surg
13 326186273...	1	p/j/s: RR RRRRR Clinical indications: Right breast inflammatory carcinoma, left axillary nodal clearance. CT - CHEST

Search History: ResultSet2

Next Result: Expand Search History Get Resultset Analytics

Prev Result: Clear Search History

Get Database Statistics

Get File Analytics

Annotation Types

Lexicon Annotation

Abbreviation

Acronym

Annotation Instances

SNOMED Display Filter

KO Display Filter

Result Page Size: 25

Load All Results

Load More Results

View in Full Screen

FIG. 7

**FRAMEWORKS AND METHODOLOGIES
FOR ENABLING SEARCHING AND/OR
CATEGORISATION OF DIGITISED
INFORMATION, INCLUDING CLINICAL
REPORT DATA**

FIELD OF THE INVENTION

[0001] The present invention relates to frameworks and methodologies for enabling categorisation and/or searching of digitised information, including clinical report data. Embodiments of the invention have been particularly developed to assist categorisation of digitised information, such as clinical report data, in a streamlined manner based on a pre-existing set of classification codes. This, in some embodiments, enables the discovery and extraction of meaningful patterns from unstructured clinical reports. Further embodiments of the invention have been particularly developed to assist in the discovery and extraction of meaningful patterns from an unstructured set of digitised information such as unstructured clinical reports. While some embodiments will be described herein with particular reference to those applications, it will be appreciated that the invention is not limited to such a field of use, and is applicable in broader contexts.

BACKGROUND

[0002] Any discussion of the background art throughout the specification should in no way be considered as an admission that such art is widely known or forms part of common general knowledge in the field.

[0003] A huge quantity of clinical record data is produced each and every day. Even for a relatively small medical facility, the task of categorising and organising clinical record data can be unwieldy. Furthermore, due to complexities of clinical language, text or keyword based searching of electronic records is limited in its efficacy.

SUMMARY OF THE INVENTION

[0004] It is an object of the present invention to overcome or ameliorate at least one of the disadvantages of the prior art, or to provide a useful alternative.

[0005] Although embodiments have been described and claimed by reference to application in fields where clinical report data is used, it should be appreciated that underlying inventions are necessarily rooted in computer technology, as opposed to abstract ideas and/or business innovations.

[0006] One embodiment provides a computer implemented method for enabling searching of clinical records, the method including:

[0007] A computer implemented method for enabling categorisation of clinical records, the method including:

[0008] identifying a set of user-generated text contained in a user interface object in a user interface environment;

[0009] processing the set of user-generated text based on a natural language processing engine, wherein the natural language processing engine leverages a clinical terminology knowledge base;

[0010] based on the processing, identifying a set of one or more suggested classification codes, wherein the classification codes are defined by a pre-existing clinical code structure;

[0011] causing display of a user interface component in the user interface environment, wherein the user interface

component: (i) displays the identified set of one or more suggested classification codes; and (ii) enables a user to selectively confirm/reject one or more of the set of identified codes;

[0012] receiving input representative of the one or more of the set of identified codes selectively confirmed by the user; and

[0013] providing a signal thereby to cause association of a clinical record including the set of user-generated text with the set of identified codes selectively confirmed by the user.

[0014] One embodiment provides a method wherein the user interface component is additionally configured to provide a search interface that enables a user to identify and selectively confirm one or more further codes, wherein the a signal causes association of a clinical record including the set of user-generated text with the set of identified codes selectively confirmed by the user and the one or more further codes.

[0015] One embodiment provides a method wherein, in response to an event whereby a user identifies and selectively confirms one or more further codes, a trigger is initiated thereby to cause updating of the natural language processing engine and/or clinical terminology knowledge base.

[0016] One embodiment provides a method including causing rendering of a graphical output that associates each identified suggested clinical classification code with a portion of the user-generated text.

[0017] One embodiment provides a method including causing rendering of a graphical output that associates each identified suggested clinical classification code with additional clinical code context.

[0018] One embodiment provides a method wherein the identification of suggested clinical classification codes is configured to identify a disease stage.

[0019] One embodiment provides a method wherein the user interface component is provided by a software application separate to a further software application in which the set of user generated text is contained.

[0020] One embodiment provides a computer implemented method for enabling searching of clinical records, the method including:

[0021] maintaining access to a database that includes data representative of clinical records, wherein each clinical record is pre-processed based upon a natural language processing engine thereby to associate the record with one or more clinical codes, wherein the clinical codes are defined in a clinical code hierarchical structure;

[0022] receiving a query from an interface that is configured enable user input of query parameters, wherein the query parameters include text-based concept-defining data;

[0023] processing the text-based concept-defining data based on a predefined protocol thereby to identify one or more clinical codes associated the concept-defining data; and

[0024] identifying one or more clinical records that are associated with the same clinical codes as the concept-defining data.

[0025] One embodiment provides a computer implemented method including additionally associating with the concept-defining data one or more further clinical codes, wherein the one or more further clinical codes are sub-tree codes to the identified one or more clinical codes in the hierarchical structure.

[0026] One embodiment provides a computer implemented method wherein the interface is configured to enable the user to select whether or not to perform the process of additionally associating with the concept-defining data one or more further clinical codes, wherein the one or more further clinical codes are sub-tree codes to the identified one or more clinical codes in the hierarchical structure.

[0027] One embodiment provides a computer implemented method wherein there is a plurality of clinical code hierarchical structures.

[0028] One embodiment provides a computer implemented method wherein the interface is configured to enable the user to select one or more of the plurality of clinical code hierarchical structures for a given query.

[0029] One embodiment provides a computer program product for performing a method as described herein.

[0030] One embodiment provides a non-transitory carrier medium for carrying computer executable code that, when executed on a processor, causes the processor to perform a method as described herein.

[0031] One embodiment provides a system configured for performing a method as described herein.

[0032] Reference throughout this specification to “one embodiment”, “some embodiments” or “an embodiment” means that a particular feature, structure or characteristic described in connection with the embodiment is included in at least one embodiment of the present invention. Thus, appearances of the phrases “in one embodiment”, “in some embodiments” or “in an embodiment” in various places throughout this specification are not necessarily all referring to the same embodiment, but may. Furthermore, the particular features, structures or characteristics may be combined in any suitable manner, as would be apparent to one of ordinary skill in the art from this disclosure, in one or more embodiments.

[0033] As used herein, unless otherwise specified the use of the ordinal adjectives “first”, “second”, “third”, etc., to describe a common object, merely indicate that different instances of like objects are being referred to, and are not intended to imply that the objects so described must be in a given sequence, either temporally, spatially, in ranking, or in any other manner.

[0034] In the claims below and the description herein, any one of the terms comprising, comprised of or which comprises is an open term that means including at least the elements/features that follow, but not excluding others. Thus, the term comprising, when used in the claims, should not be interpreted as being limitative to the means or elements or steps listed thereafter. For example, the scope of the expression a device comprising A and B should not be limited to devices consisting only of elements A and B. Any one of the terms including or which includes or that includes as used herein is also an open term that also means including at least the elements/features that follow the term, but not excluding others. Thus, including is synonymous with and means comprising.

[0035] As used herein, the term “exemplary” is used in the sense of providing examples, as opposed to indicating quality. That is, an “exemplary embodiment” is an embodiment provided as an example, as opposed to necessarily being an embodiment of exemplary quality.

BRIEF DESCRIPTION OF THE DRAWINGS

[0036] Embodiments of the invention will now be described, by way of example only, with reference to the accompanying drawings in which:

[0037] FIG. 1 schematically illustrates a framework according to one embodiment.

[0038] FIG. 2 illustrates a method according to one embodiment.

[0039] FIG. 3 illustrates an exemplary client-server framework.

[0040] FIG. 4 illustrates a search interface according to one embodiment.

[0041] FIG. 5 schematically illustrates a framework according to one embodiment.

[0042] FIG. 6 illustrates a method according to one embodiment.

[0043] FIG. 7 illustrates a search interface according to one embodiment.

DETAILED DESCRIPTION

[0044] Described herein are frameworks and methodologies for enabling categorisation of clinical report data. Embodiments of the invention have been particularly developed to assist categorisation of digitised information, such as clinical report data, in a streamlined manner based on a pre-existing set of classification codes. This, in some embodiments, enables the discovery and extraction of meaningful patterns from unstructured clinical reports. Further embodiments of the invention have been particularly developed to assist in the discovery and extraction of meaningful patterns from an unstructured set of digitised information such as unstructured clinical reports. While some embodiments will be described herein with particular reference to those applications, it will be appreciated that the invention is not limited to such a field of use, and is applicable in broader contexts.

Abstraction Technology

[0045] In the context of clinical report management, there is a technical problem to be solved in the context of codifying reports based on predefined schema (such as classification codes based on SNOMED CT, ICD, LOINC, and ICD-O3). Typically this is performed by a report Author, or (perhaps more commonly) by an assistant who reviews report data. A partial solution is to provide a search interface which enables an author (or reviewer) to perform searches based on subjective analysis of report content (either at the time of reading or later on). However, this is time-consuming and fraught with error potential. Ideally, computer technology would be applied to enable automated coding, but subjective factors (such as author style and terminology variances) render that somewhat unrealistic. The technology described below provides a technical solution in the form of a particular computer interface which enables a streamlined assistive process towards report coding.

[0046] In some embodiments, the technology includes a software component described herein as an “abstractor component”. This component is configured to provide automation of abstracting and coding of digitised information. For the purposes clinical text (for example clinical reports, which may include notes and the like). This makes use of a natural language processing (NLP) engine (for example a text analytics engine) that operates in conjunction with (and

leverages) a clinical terminology knowledge base. This clinical terminology base provides an expandable intimate knowledge of clinical terminology above, beyond a general purpose natural language processing engine.

[0047] The abstractor, in combination with NLP engine and clinical terminology knowledge base, enables concept-based abstracting of clinical records. This allows technology to operate in an objective manner which effectively ignores subjective variance factors such as idiosyncrasies of clinicians' writing styles, and language usage to accurately code records. The abstractor component is in some embodiments configured to operate within an existing software environment, directly interact with an existing software environment (for example via an API), or indirectly interact with an existing software environment (for example via operating system level functionalities).

[0048] Various embodiments provide computer implemented methods for enabling categorisation of clinical records. Such methods include identifying a set of user-generated text contained in a user interface object in a user interface environment. For example, in some cases this is: text contained in an active object, such as a specific data field rendered by a given software application; text contained in an active window in third party software; text contained in a referenced file, or the like. In some cases the set of user-generated text includes all text in the object; in other cases it is a reduced selection (for example a user-designated selection, which is for instance made via a highlighting function).

[0049] The methods then include processing the set of user-generated text based on a natural language processing engine, being a natural language processing engine leveraging a clinical terminology knowledge base. This knowledge base provides rules and/or relationships to assist in the processing of known medical/clinical terms and expressions based on: known word morphology; sentence parsing; spelling correction; and acronym/abbreviation resolution examples. The knowledge based is preferably updated over time (for example via a lexicon management system) thereby to enable identification of and accounting for additional clinician idiosyncratic authoring styles, and the like.

[0050] Based on the processing, a set of one or more suggested classification codes is identified. The classification codes are defined by a pre-existing clinical code structure, such as SNOMED CT, ICD, LOINC, and ICD-O3. In some embodiments multiple code structures are used, and in some cases a user is enabled to select one or more of a set of available code structures (for example via check-boxes or the like). This selection may occur following identification of codes via a user interface component, as discussed below.

[0051] Following identification of a set of one or more codes, the methods include causing display of a particular functionally relevant user interface component in the user interface environment. This may in some embodiments be a pop-up window, which is triggered by an external software component. The user interface component is configured to: (i) display the identified set of one or more suggested classification codes; and (ii) enable a user to selectively confirm/reject one or more of the set of identified codes. For example, in a preferred embodiment the user interface component displays, for each identified code, some or all of the following aspects of information:

[0052] A section of the code structure to which an identified code belongs.

[0053] A disease class associated with an identified code.

[0054] A disease stage associated with an identified code. For example, in one embodiment the natural language processing engine is configured to make inferences about the data, for example; inferring a stage of cancer diagnosis.

[0055] A disease behaviour associated with an identified code.

[0056] A code-structure-specific description associated with an identified code.

[0057] A portion of text that triggered the identification of an identified code.

[0058] A portion of contextual text including a sub-portion of text that triggered the identification of an identified code. In some embodiments an author is enabled to modify this portion of text in the interface, optionally causing re-processing and re-population of user interface data, and trigger automated updating of the source clinical report text to account for the modification.

[0059] This content is preferably displayed in a user-friendly tabulation, which is in some embodiments able to subsequently be stored (or linked to) a file or data set defining a clinical record defined from the source text.

[0060] The methods also include receiving input representative of the one or more of the set of identified codes selectively confirmed by the user (for example by confirm/reject check boxes provided via the user interface). This triggers the providing of a signal thereby to cause association of a clinical record including the set of user-generated text with the set of identified codes selectively confirmed by the user. For example, the association may be by way of added textual information, metadata, database level relationships, and so on.

[0061] In some embodiments, the user interface component is additionally configured to provide a search interface that enables a user to identify and selectively confirm one or more further codes (which are also associated with the record). For example, an author may notice one or more code descriptors (such as particular diseases) missing, and search for those. This may, in some cases, encourage the author to modify the text, thereby to include material that had been inadvertently omitted (or to correct spelling mistakes and the like).

[0062] In some cases, in response to an event whereby a user identifies and selectively confirms one or more further codes, a trigger is initiated thereby to cause updating of the natural language processing engine and/or clinical terminology knowledge base. This may be a manual or automated process, whereby the text is analysed thereby to determine whether there is content that should have caused the automated identification of a manually added clinical code, this resulting in updating of the processing engine and/or clinical terminology knowledge base such that the relevant code is identified in a corresponding situation in future.

[0063] It should be appreciated that, in preferred embodiments, the textual language processing engine is configured to identify discrete portions of the text, and, identify a set of one or more suggested classification codes for each discrete portion (as opposed to processing the document as a single whole textual entity). In some embodiments association of the clinical record with the set of identified codes includes defining data that enables association of each discrete por-

tion with its respective classification codes (such that a user is able to identify a particular sub-portion of text that resulted in a given code association).

[0064] FIG. 1 illustrates a framework according to one embodiment. In this example, various user interface objects are schematically illustrated as being rendered by an exemplary clinical record system **100**. More specifically, these objects are illustrated as being rendered in a common software application, however in further embodiments they are rendered via multiple different software applications (for example using plug-ins and the like).

[0065] A user authors text for inclusion in (or defining) a clinical record in object **101**. Then, by way of a trigger object (such as a “submit” or “code” button), the text data is provided to a record analysis subsystem **120**. Subsystem **120** includes natural language processing logic **121** (which leverages an expandable predefined clinical terminology knowledge base), and a repository of clinical code data (which includes textual description data associated with each of a plurality of clinical code identifiers for one or more clinical coding schema). A record processing engine is configured to coordinate the processing of text data retrieved from system **100**, and provide data representative of identified codes for rendering in a user interface object **103**.

[0066] User interface object **103** provides data representative of identified codes, in this case including, for each identified code, a code descriptor and one or more artefacts of code context data (for example retrieved from code data **123** and/or extracted from the source text from object **101**). It also provides confirm/reject controls (which may take the form of checkboxes), thereby to enable user selective confirm/reject command inputs.

[0067] A button **103** provides a user with access to a search interface, which enables text-based searching of data **123** thereby to identify one or more further clinical codes manually. In some cases each interface also leverages logic **121** and engine **122** thereby to expand a search string to assist in identification of codes whilst taking into account morphology, spelling, acronyms, and many other language variables.

[0068] In the example of FIG. 1, records are stored in a database **110** via subsystem **120**. However, in another example, interface **103** is configured to modify text in object **101** (thereby to append clinical codes), or otherwise modify operation of a pre-existing software component that causes recording of a clinical record based on text in object **101** (for example by metadata or database-level association).

[0069] FIG. 2 illustrates a method according to one embodiment. A user defines clinical record text at **201**, and then a trigger event is caused at **202** (for example subject to a user interaction following completion of text authoring). Natural language processing of the text occurs at **203**, resulting in identification of one or more predicted/suggested clinical codes at **204**. The clinical codes are presented by way of a user interface at **205**, which enables user review and selective confirmation. Optionally, at **206**, a user accesses a search interface to manually identify and add one or more further codes. User confirmation of codes is received at **207**, and the codes and record stored at **208**. This may include adding the codes to the text content of a clinical record, and storing that clinical record in a database.

[0070] It will be appreciated that the technology described above is not limited to the particular practical implementation environment of clinical record management. For

instance, the computer technology in question may be implemented in a variety of other situations where user-defined text requires coding according to a predefined schema.

Concept Search Functionalities

[0071] Described below are frameworks and methodologies for enabling searching of clinical report data. Embodiments have been particularly developed to assist in the discovery and extraction of meaningful patterns from unstructured clinical reports. While some embodiments will be described herein with particular reference to that application, it will be appreciated that the invention is not limited to such a field of use, and is applicable in broader contexts.

[0072] Various embodiments provide a computer implemented methods for enabling searching of clinical records. For example, this in some cases includes providing a search engine that is configured to access clinical record data provided through one or more clinical systems (or other record generation means). The methods include maintaining access to a database that includes data representative of clinical records. For example, this may be a locally or remotely hosted database, and in some embodiments includes multiple distributed databases.

[0073] Each clinical record is pre-processed based upon a natural language processing engine, thereby to associate the record with one or more clinical codes 9 for example via technology described further above). These clinical codes are, at least in some examples, defined in a clinical code hierarchical structure. An example of a clinical code hierarchical structure used in some embodiments is SNOMED CT. Other code structures include ICD, LOINC, and ICD-O3. In some embodiments there is a plurality of clinical code hierarchical structures, and a user is enabled to select one or more of these to be used for the purpose of a given query.

[0074] The natural language processing engine preferably accounts factors such as word morphology, sentence parsing, spelling correction, synonym, metonym, acronym and abbreviation resolution, document structure thereby to enable mapping of records to codes in spite of a wide range of language variations that may be observed in clinical records. The natural language processing engine is configured to be updated (and re-run over some or all of the records) over time, thereby to improve its effectiveness.

[0075] A search interface is provided, for example in a manner enabling rendering of a query-defining interface at a client terminal. Such an interface configured to enable user input of query parameters. These query parameters include text-based concept-defining data, for example “breast cancer” or the like.

[0076] A query processing engine is configured to process the text-based concept-defining data based on a predefined protocol thereby to identify one or more clinical codes associated the concept-defining data. For example, factors such as word morphology, sentence parsing, spelling correction, acronym and abbreviation resolution are handled by such a protocol, which is configured to map various instances of concept-defining data to codes, even if there is no direct textual relationship between a code descriptor and the concept-defining data itself. In some embodiments the concept-defining data is defined by a text-based set of search terms (for example a search string), which only one representation of the semantics of a concept, being a concept that is in effect searched as described herein. For example, the

concept-defining data provides a single representation of a concept that has an non-enumerable set of representations [0077] In overview, the query processing engine is configured to identify one or more clinical records that are associated with the same clinical codes as a concept represented by the instance of the concept-defining data. In this regard, there is a dual-level expansion process:

[0078] Firstly, expansion in relation to the text of a clinical record thereby to associate it with codes based on natural language processing; and

[0079] Secondly, expansion at the concept-defining data end to identify relevant codes.

[0080] By way of example, consider the following:

[0081] A report may includes a term “Term A”.

[0082] Based on enhanced NLP-based processing, that term is determined to match “Term B” (for example they are determined to be equivalents).

[0083] Term B is associated with Code X.

[0084] A query includes Term C.

[0085] Term C is processed also found to be associated with Code X. This may include processing that associates Term C with Term B term via enhanced NLP-based processing, before arriving at the association with Code X.

[0086] Examples of Term A, B and C might be “Lentigo Maligna”, “melanoma” and “skin cancer”.

[0087] The above example, a search using Term C is executed to identify the record containing Term A, on the basis that they both match Term B of Code X. This provides a significant technical advantage over conventional searching technology, on the basis that term C and Term A might not be directly relatable.

[0088] In some embodiments the method additionally includes associating with the concept-defining data one or more further clinical codes, wherein the one or more further clinical codes are sub-tree codes to the identified one or more clinical codes in the hierarchical structure. For example, the search interface is configured to enable the user to select whether or not to perform the process of additionally associating with the concept-defining data one or more further clinical codes, wherein the one or more further clinical codes are sub-tree codes to the identified one or more clinical codes in the hierarchical structure. An example in the context of medical terminology is “breast cancer” which has the sub-type “infiltrating ductal carcinoma”.

[0089] FIG. 1 illustrates a framework according to one embodiment. A plurality of clinical records defining systems 500 include an exemplary client terminal 500'. These are each optionally used to define clinical record data. This data is provided to a central clinical record database 510. Systems 500 may include various software platforms used to generate clinical report content. A record analysis subsystem includes a record processing engine 532, which is configured to maintain access to database 510, and process records therein based on natural language processing logic 531. More specifically, this processing enables each record to be associated with one or more clinical codes, based on clinical code data 533. Record processing engine 532 is preferably configured to process newly added records, and to re-process existing records following updates to logic 531 and/or data 533.

[0090] A search engine subsystem 530 is configured to process queries submitted from client terminals 550, including an exemplary client terminal 550'. The client terminals

in some embodiments provide search interfaces via proprietary software, or in some cases via web-browser delivered interfaces.

[0091] FIG. 2 illustrates exemplary methods 600 and 610. Method 600 is a method for natural language processing of clinical record data thereby to enable association with codes (such as SNOMED CT codes), whereas method 610 is a query processing method.

[0092] In relation to method 600, functional block 601 represents a process including identifying a clinical record for processing. For example this may be a newly submitted record, or a record queued for periodic re-processing (for example based on updated processing logic). Functional block 602 represents a process including application of natural language processing logic, which enables association of the record to codes (such as SNOMED CT codes), whilst resolving amongst other things word morphology, sentence parsing, spelling correction, acronym and abbreviation resolution. Code associations are stored in the database at 604.

[0093] In relation to method 610, functional block 611 includes receiving a query. The query is processed thereby to identify codes (such as SNOMED CT codes) at 612, again resolving issues such as word morphology, sentence parsing, spelling correction, acronym and abbreviation resolution. This enables translation between terms used by a query submitting user, and terms actually present in code descriptors. As represented by functional block 613, is the user a selected an option for an “expanded” query, then sub-tree codes are also determined. Functional block 614 represents a process including identifying those records associated with the same codes (and, where relevant, sub-tree codes).

[0094] FIG. 4 illustrates an exemplary query interface. In this example, a user inputs a concept-defining data in a “terminology code or description” field, and may also input excluded concepts. The user also selects codes to be considered (in this case being ICD-10 and/or SNOMED), and optionally selects whether a sub-tree search is to be performed. The user also selects whether “any” or “all” type search logic is to be applied. Various options for search refining/combining are also provided. In the illustrated embodiment, search results are displayed along with a window that displays content of a selected search result.

[0095] It will be appreciated that technology described herein is readily configured to assist in discovering and extracting meaningful patterns (analytics) in patient data within unstructured clinical reports. These are able to provide fast search speed and precision extraction, using combinations of multiple search algorithms and language processing techniques to maximise the coverage and reliability of results, may co-ordinate with any search engine to multiply the value of locally developed expertise, and can retrieve reports in tandem with structured retrieval queries to collect variable data types into the one location.

[0096] The use of a concept-based search mechanism enables understanding of the idiosyncrasies of clinicians’ writing styles and language usage, thereby to accurately collect desired data. In some embodiments the technology is configured to extract the appropriate content for either display in a user-friendly tabulation or export to a CSV file for subsequent processing by combining a variety of facts in a report. Furthermore, it is in some embodiments set up to produce tailored reports to support each organisation’s specific needs.

[0097] In some embodiments efficacy is improved by updating natural language processing logic by tuning it to local content in a process of training it to understand the writing vagaries of local authors. Natural language engine resolving amongst other things word morphology, sentence parsing, spelling correction, acronym and abbreviation resolution.

[0098] Concept processing may include identifying any text equivalent in a concept-defining data to ICD, LOINC, O3 and SNOMED CT codes (and in some cases proprietary dictionaries).

[0099] Extractions may include disease, disorders, procedures, with special attention to cancer details for primaries, metastases and nodes. Document analytics may be performed, including classification of documents from their contents such as for tumour stream or report purpose.

[0100] In some embodiments external reuse of extractions & inferences is provided, for example including buttons to export extracted content to CVS files for delivery to researchers or search collections.

[0101] In some embodiments workflow ease is assisted by providing a search history, batch processing and results export so as to assemble a workflow and operate it entirely automatically thus supporting periodic reporting.

[0102] In some embodiments a mechanism is provided to capture content that is not processed initially can be used to lodge that content into the knowledge repository so it is not missed again.

[0103] Some embodiments provide automatic despatch of files to registered users when a record of interest is created in the system.

[0104] Embodiments may, by way of example, be applied in the following environments:

[0105] Patient Studies (for example to recruit a common cohort or phenotypes for clinical trials; to target particular disease classes for research; to alert staff when a pertinent record is created in the storage system).

[0106] Auditing Clinical Records (for example to check the accuracy of clinical coding across disease classes; to semi-automate codification for billing; to identify and understand the case mix of a particular organisation).

[0107] Research (for example to answer ad-hoc questions about the distributions of disease morbidities).

[0108] Assessing Training & Report Writing (for example to investigate report reliability and readability; to provide feedback to staff on the nature of their written composition; to tutor and evaluate the work of trainees; to compile activity reports for professional associations for registration; to investigate more effective systematic methods for preparing reports).

[0109] It will be appreciated that these technologies overcome technical problems associated with automated searching and/or processing of authored report data, where that data include idiosyncratic and/or subjectively defined elements.

Example Client-Server Framework

[0110] In some embodiments, methods and functionalities considered herein are implemented by way of a client-server framework, as illustrated in FIG. 3. In overview, a web server 302 provides a web interface 303. This web interface is accessed by the parties by way of client terminals 304. In overview, users access interface 303 over the Internet by

way of client terminals 304, which in various embodiments include the likes of personal computers, PDAs, cellular telephones, gaming consoles, and other Internet enabled devices.

[0111] Server 303 includes a processor 305 coupled to a memory module 306 and a communications interface 307, such as an Internet connection, modem, Ethernet port, wireless network card, serial port, or the like. In other embodiments distributed resources are used. For example, in one embodiment server 302 includes a plurality of distributed servers having respective storage, processing and communications resources. Memory module 306 includes software instructions 308, which are executable on processor 305.

[0112] Server 302 is coupled to a database 310. In further embodiments the database leverages memory module 306.

[0113] In some embodiments web interface 303 includes a website. The term “website” should be read broadly to cover substantially any source of information accessible over the Internet or another communications network (such as WAN, LAN or WLAN) via a browser application running on a client terminal. In some embodiments, a website is a source of information made available by a server and accessible over the Internet by a web-browser application running on a client terminal. The web-browser application downloads code, such as HTML code, from the server. This code is executable through the web-browser on the client terminal for providing a graphical and often interactive representation of the website on the client terminal. By way of the web-browser application, a user of the client terminal is able to navigate between and throughout various web pages provided by the website, and access various functionalities that are provided.

[0114] Although some embodiments make use of a website/browser-based implementation, in other embodiments proprietary software methods are implemented as an alternative. For example, in such embodiments client terminals 304 maintain software instructions for a computer program product that essentially provides access to a portal via which framework 100 is accessed (for instance via an iPhone app or the like).

[0115] In general terms, each terminal 304 includes a processor 311 coupled to a memory module 313 and a communications interface 312, such as an internet connection, modem, Ethernet port, serial port, or the like. Memory module 313 includes software instructions 314, which are executable on processor 311. These software instructions allow terminal 304 to execute a software application, such as a proprietary application or web browser application and thereby render on-screen a user interface and allow communication with server 302. This user interface allows for the creation, viewing and administration of profiles, access to the internal communications interface, and various other functionalities.

[0116] Unless specifically stated otherwise, as apparent from the following discussions, it is appreciated that throughout the specification discussions utilizing terms such as “processing,” “computing,” “calculating,” “determining”, “analyzing” or the like, refer to the action and/or processes of a computer or computing system, or similar electronic computing device, that manipulate and/or transform data represented as physical, such as electronic, quantities into other data similarly represented as physical quantities.

[0117] In a similar manner, the term “processor” may refer to any device or portion of a device that processes electronic data, e.g., from registers and/or memory to transform that electronic data into other electronic data that, e.g., may be stored in registers and/or memory. A “computer” or a “computing machine” or a “computing platform” may include one or more processors.

[0118] The methodologies described herein are, in one embodiment, performable by one or more processors that accept computer-readable (also called machine-readable) code containing a set of instructions that when executed by one or more of the processors carry out at least one of the methods described herein. Any processor capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken are included. Thus, one example is a typical processing system that includes one or more processors. Each processor may include one or more of a CPU, a graphics processing unit, and a programmable DSP unit. The processing system further may include a memory subsystem including main RAM and/or a static RAM, and/or ROM. A bus subsystem may be included for communicating between the components. The processing system further may be a distributed processing system with processors coupled by a network. If the processing system requires a display, such a display may be included, e.g., a liquid crystal display (LCD) or a cathode ray tube (CRT) display. If manual data entry is required, the processing system also includes an input device such as one or more of an alphanumeric input unit such as a keyboard, a pointing control device such as a mouse, and so forth. The term memory unit as used herein, if clear from the context and unless explicitly stated otherwise, also encompasses a storage system such as a disk drive unit. The processing system in some configurations may include a sound output device, and a network interface device. The memory subsystem thus includes a computer-readable carrier medium that carries computer-readable code (e.g., software) including a set of instructions to cause performing, when executed by one or more processors, one of more of the methods described herein. Note that when the method includes several elements, e.g., several steps, no ordering of such elements is implied, unless specifically stated. The software may reside in the hard disk, or may also reside, completely or at least partially, within the RAM and/or within the processor during execution thereof by the computer system. Thus, the memory and the processor also constitute computer-readable carrier medium carrying computer-readable code.

[0119] Furthermore, a computer-readable carrier medium may form, or be included in a computer program product.

[0120] In alternative embodiments, the one or more processors operate as a standalone device or may be connected, e.g., networked to other processor(s), in a networked deployment, the one or more processors may operate in the capacity of a server or a user machine in server-user network environment, or as a peer machine in a peer-to-peer or distributed network environment. The one or more processors may form a personal computer (PC), a tablet PC, a set-top box (STB), a Personal Digital Assistant (PDA), a cellular telephone, a web appliance, a network router, switch or bridge, or any machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine.

[0121] Note that while diagrams only show a single processor and a single memory that carries the computer-

readable code, those in the art will understand that many of the components described above are included, but not explicitly shown or described in order not to obscure the inventive aspect. For example, while only a single machine is illustrated, the term “machine” shall also be taken to include any collection of machines that individually or jointly execute a set (or multiple sets) of instructions to perform any one or more of the methodologies discussed herein.

[0122] Thus, one embodiment of each of the methods described herein is in the form of a computer-readable carrier medium carrying a set of instructions, e.g., a computer program that is for execution on one or more processors, e.g., one or more processors that are part of web server arrangement. Thus, as will be appreciated by those skilled in the art, embodiments of the present invention may be embodied as a method, an apparatus such as a special purpose apparatus, an apparatus such as a data processing system, or a computer-readable carrier medium, e.g., a computer program product. The computer-readable carrier medium carries computer readable code including a set of instructions that when executed on one or more processors cause the processor or processors to implement a method. Accordingly, aspects of the present invention may take the form of a method, an entirely hardware embodiment, an entirely software embodiment or an embodiment combining software and hardware aspects. Furthermore, the present invention may take the form of carrier medium (e.g., a computer program product on a computer-readable storage medium) carrying computer-readable program code embodied in the medium.

[0123] The software may further be transmitted or received over a network via a network interface device. While the carrier medium is shown in an exemplary embodiment to be a single medium, the term “carrier medium” should be taken to include a single medium or multiple media (e.g., a centralized or distributed database, and/or associated caches and servers) that store the one or more sets of instructions. The term “carrier medium” shall also be taken to include any medium that is capable of storing, encoding or carrying a set of instructions for execution by one or more of the processors and that cause the one or more processors to perform any one or more of the methodologies of the present invention. A carrier medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical, magnetic disks, and magneto-optical disks. Volatile media includes dynamic memory, such as main memory. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise a bus subsystem. Transmission media also may also take the form of acoustic or light waves, such as those generated during radio wave and infrared data communications. For example, the term “carrier medium” shall accordingly be taken to include, but not be limited to, solid-state memories, a computer product embodied in optical and magnetic media; a medium bearing a propagated signal detectable by at least one processor of one or more processors and representing a set of instructions that, when executed, implement a method; and a transmission medium in a network bearing a propagated signal detectable by at least one processor of the one or more processors and representing the set of instructions.

[0124] It will be understood that the steps of methods discussed are performed in one embodiment by an appropriate processor (or processors) of a processing (i.e., computer) system executing instructions (computer-readable code) stored in storage. It will also be understood that the invention is not limited to any particular implementation or programming technique and that the invention may be implemented using any appropriate techniques for implementing the functionality described herein. The invention is not limited to any particular programming language or operating system.

[0125] It should be appreciated that in the above description of exemplary embodiments of the invention, various features of the invention are sometimes grouped together in a single embodiment, FIG., or description thereof for the purpose of streamlining the disclosure and aiding in the understanding of one or more of the various inventive aspects. This method of disclosure, however, is not to be interpreted as reflecting an intention that the claimed invention requires more features than are expressly recited in each claim. Rather, as the following claims reflect, inventive aspects lie in less than all features of a single foregoing disclosed embodiment. Thus, the claims following the Detailed Description are hereby expressly incorporated into this Detailed Description, with each claim standing on its own as a separate embodiment of this invention.

[0126] Furthermore, while some embodiments described herein include some but not other features included in other embodiments, combinations of features of different embodiments are meant to be within the scope of the invention, and form different embodiments, as would be understood by those skilled in the art. For example, in the following claims, any of the claimed embodiments can be used in any combination.

[0127] Furthermore, some of the embodiments are described herein as a method or combination of elements of a method that can be implemented by a processor of a computer system or by other means of carrying out the function. Thus, a processor with the necessary instructions for carrying out such a method or element of a method forms a means for carrying out the method or element of a method. Furthermore, an element described herein of an apparatus embodiment is an example of a means for carrying out the function performed by the element for the purpose of carrying out the invention.

[0128] In the description provided herein, numerous specific details are set forth. However, it is understood that embodiments of the invention may be practiced without these specific details. In other instances, well-known methods, structures and techniques have not been shown in detail in order not to obscure an understanding of this description.

[0129] Similarly, it is to be noticed that the term coupled, when used in the claims, should not be interpreted as being limited to direct connections only. The terms “coupled” and “connected,” along with their derivatives, may be used. It should be understood that these terms are not intended as synonyms for each other. Thus, the scope of the expression a device A coupled to a device B should not be limited to devices or systems wherein an output of device A is directly connected to an input of device B. It means that there exists a path between an output of A and an input of B which may be a path including other devices or means. “Coupled” may mean that two or more elements are either in direct physical

or electrical contact, or that two or more elements are not in direct contact with each other but yet still co-operate or interact with each other.

[0130] Thus, while there has been described what are believed to be the preferred embodiments of the invention, those skilled in the art will recognize that other and further modifications may be made thereto without departing from the spirit of the invention, and it is intended to claim all such changes and modifications as falling within the scope of the invention. For example, any formulas given above are merely representative of procedures that may be used. Functionality may be added or deleted from the block diagrams and operations may be interchanged among functional blocks. Steps may be added or deleted to methods described within the scope of the present invention.

What is claimed is:

1-9. (canceled)

10. A computer implemented method for enabling searching of clinical records the method including:

maintaining access to a database that includes data representative of clinical records, wherein each clinical record is pre-processed based upon a natural language processing engine thereby to associate the record with one or more clinical codes, wherein the clinical codes are defined in a clinical code hierarchical structure;

receiving a query from an interface that is configured enable user input of query parameters, wherein the query parameters include text-based concept-defining data;

processing the text-based concept-defining data based on a predefined protocol thereby to identify one or more clinical codes associated with the concept-defining data; and

identifying one or more clinical records that are associated with the same clinical codes as the concept-defining data.

11. A method according to claim **10** including additionally associating with the concept-defining data one or more further clinical codes, wherein the one or more further clinical codes are sub-tree codes to the identified one or more clinical codes in the hierarchical structure.

12. A method according to claim **11** wherein the interface is configured to enable the user to select whether or not to perform the process of additionally associating with the concept-defining data one or more further clinical codes, wherein the one or more further clinical codes are sub-tree codes to the identified one or more clinical codes in the hierarchical structure.

13. A method according to claim **10**, wherein there is a plurality of clinical code hierarchical structures.

14. A method according to claim **13** wherein the interface is configured to enable the user to select one or more of the plurality of clinical code hierarchical structures for a given query.

15. A method according to claim **10** wherein the concept-defining data provides a single representation of a concept that has a non-enumerable set of representations.

16. A computer implemented method for enabling categorisation and searching of clinical records, the method including:

identifying a set of user-generated text contained a user interface object in a user interface environment;

processing the set of user-generated text based on a natural language processing engine, wherein the natural language processing engine leverages a clinical terminology knowledge base;

based on the processing, identifying a set of one or more suggested classification codes, wherein the classification codes are defined by a pre-existing clinical code structure;

causing display of a user interface component in the user interface environment, wherein the user interface component: (i) displays the identified set of one or more suggested classification codes; and (ii) enables a user to selectively confirm/reject one or more of the set of identified codes;

receiving input representative of the one or more of the set of identified codes selectively confirmed by the user;

providing a signal thereby to cause association of a clinical record including the set of user-generated text with the set of identified codes selectively confirmed by the user;

maintaining access to a database that includes data representative of clinical records, wherein each clinical record is pre-processed based upon a natural language

processing engine thereby to associate the record with one or more clinical codes, wherein the clinical codes are defined in a clinical code hierarchical structure;

receiving a query from an interface that is configured enable user input of query parameters, wherein the query parameters include text-based concept-defining data;

processing the text-based concept-defining data based on a predefined protocol thereby to identify one or more clinical codes associated with the concept-defining data; and

identifying one or more clinical records that are associated with the same clinical codes as the concept-defining data.

17. A computer system configured to perform a method according to claim **10**.

18. A computer program configured to perform a method according to claim **10**.

19. A non-transitory carrier medium carrying computer executable code that, when executed on a processor, causes the processor to perform a method according to claim **10**.

* * * * *