



US 20220310260A1

(19) **United States**

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

(10) **Pub. No.: US 2022/0310260 A1**

(43) **Pub. Date: Sep. 29, 2022**

(54) **SYSTEM AND METHODS FOR
KNOWLEDGE REPRESENTATION AND
REASONING IN CLINICAL PROCEDURES**

Publication Classification

(51) **Int. Cl.**
G16H 50/20 (2006.01)
(52) **U.S. Cl.**
CPC **G16H 50/20** (2018.01)

(71) Applicant: **Siemens Healthcare GmbH**, Erlangen
(DE)

(72) Inventors: **Poikavila Ullaskrishnan**, Lebanon, NH
(US); **Tiziano Passerini**, Plainsboro, NJ
(US); **Puneet Sharma**, Princeton
Junction, NJ (US); **Paul Klein**,
Princeton, NJ (US); **Teodora-Vanessa
Liliac**, Brasov (RO); **Larisa Micu**,
Brasov (RO)

(57) **ABSTRACT**

A medical knowledge base in a digital, clinical system is upgraded. A storage with a knowledge base, being a SNOMED knowledge base, is provided in a web ontology format. Procedural data, representing clinical procedures for evaluation of a patient's health state, is received. The received procedural data is mapped in a set of SNOMED expressions. The SNOMED expressions are converted into statements in the web ontology format. The SNOMED knowledge base is upgraded with the received procedural data by adding the statements in the SNOMED knowledge base for providing a processable file with an upgraded version of the SNOMED knowledge base.

(21) Appl. No.: **17/655,001**

(22) Filed: **Mar. 16, 2022**

(30) **Foreign Application Priority Data**

Mar. 29, 2021 (EP) 21146514.4




FIG 1




FIG 2




FIG 3




FIG 4




FIG 5

File Edit View Reasoner Tools Refactor Window Help			
< > version (http://snomed_info/sct/900000000000207008/version/20200818)			
Active ontology x Entities x Individuals by class x DL Query x OntoGraf x SPARQL Query x			
Individuals	≡ ◇ P1 - http://snomed_info/sct/900000000000207008/P1		
Datatypes			
Annotation pro.			
Data properties			
Object propert.			
Classes			
Indivi			
◇+ ✕			
◆ P1			
◇ P2			
◇ P3			
Description: P1		Property assertions: P1	
Types + ○ 'Patient (person)'		Object property assertions +	
Same Individual As +		Data property assertions +	
Different Individuals +		■ 'Aortic valve area Maximum Range (observable entity)' "1.05" ^ ^ xsd:double	
		■ 'Peak arterial velocity Maximum Range (observable entity)' "4.305" ^ ^ xsd:double	
		■ 'Peak arterial velocity Minimum Range (observable entity)' "3.895" ^ ^ xsd:double	
		■ 'Cardiovascular pressure gradient Maximum Range (observable entity)' "52.5" ^ ^ xsd:double	
		■ 'Aortic valve area Minimum Range (observable entity)' "0.95" ^ ^ xsd:double	
		■ 'Cardiovascular pressure gradient Minimum Range (observable entity)' "47.5" ^ ^ xsd:double	
		Negative object property assertions +	
		Negative data property assertions +	

File Edit View Reasoner Tools Refactor Window Help			
< > version (http://snomed_info/sct/900000000000207008/version/20200818)			
Active ontology x Entities x Individuals by class x DL Query x OntoGraf x SPARQL Query x			
Individuals	≡ ◇ P1 - http://snomed_info/sct/900000000000207008/P1		
Datatypes			
Annotation pro.			
Data properties			
Object propert.			
Classes			
Indivi			
◇+ ✕			
◆ P1			
◇ P2			
◇ P3			
Description: P1		Property assertions: P1	
Types + ○ 'Patient (person)' ● 'Severe stenosis of aortic valve (disorder)'		Object property assertions +	
Same Individual As +		Data property assertions +	
Different Individuals +		■ 'Aortic valve area Maximum Range (observable entity)' "1.05" ^ ^ xsd:double	
		■ 'Peak arterial velocity Maximum Range (observable entity)' "4.305" ^ ^ xsd:double	
		■ 'Peak arterial velocity Minimum Range (observable entity)' "3.895" ^ ^ xsd:double	
		■ 'Cardiovascular pressure gradient Maximum Range (observable entity)' "52.5" ^ ^ xsd:double	
		■ 'Aortic valve area Minimum Range (observable entity)' "0.95" ^ ^ xsd:double	
		■ 'Cardiovascular pressure gradient Minimum Range (observable entity)' "47.5" ^ ^ xsd:double	
		■ 'Aortic valve area (observable entity)' "1.05" ^ ^ xsd:double	
		■ 'Aortic valve area (observable entity)' "0.95" ^ ^ xsd:double	
		■ 'Cardiovascular pressure gradient (observable entity)' "47.5" ^ ^ xsd:double	
		■ 'Cardiovascular pressure gradient (observable entity)' "52.5" ^ ^ xsd:double	
		■ 'Peak arterial velocity (observable entity)' "4.305" ^ ^ xsd:double	
		■ 'Peak arterial velocity (observable entity)' "3.895" ^ ^ xsd:double	
		Negative object property assertions +	
		Negative data property assertions +	

FIG 6

File Edit View Reasoner Tools Refactor Window Help

< >

version (http://snomed_info/sct/90000000000207008/version/20200818)

Active ontology x Entities x Individuals by class x DL Query x OntoGraf x SPARQL Query x

≡

◇ P2 - http://snomed_info/sct/90000000000207008/P1

Individuals

Datatypes

Annotation pro.

Data properties

Object propert.

Classes

Indivi

◇ + ✕

◇ P1

◇ P2

◇ P3

Description: P2

Types

○ 'Patient (person)'

⊖ 'Moderate stenosis of aortic valve (disorder)'

⊖ 'Well adult (finding)'

Same Individual As

Different Individuals

Property assertions: P2

Object property assertions

Data property assertions

'Peak arterial velocity Maximum Range (observable entity)' "3.78" ^ ^ xsd:double

'Aortic valve area Maximum Range (observable entity)' "1.05" ^ ^ xsd:double

'Peak arterial velocity Minimum Range (observable entity)' "3.42" ^ ^ xsd:double

'Cardiovascular pressure gradient Minimum Range (observable entity)' "35.15" ^ ^ xsd:double

'Cardiovascular pressure gradient Maximum Range (observable entity)' "38.85" ^ ^ xsd:double

'Aortic valve area Minimum Range (observable entity)' "0.95" ^ ^ xsd:double

'Aortic valve area (observable entity)' "1.05" ^ ^ xsd:double

'Aortic valve area (observable entity)' "0.95" ^ ^ xsd:double

'Cardiovascular pressure gradient (observable entity)' "35.15" ^ ^ xsd:double

'Cardiovascular pressure gradient (observable entity)' "38.85" ^ ^ xsd:double

'Peak arterial velocity (observable entity)' "3.42" ^ ^ xsd:double

'Peak arterial velocity (observable entity)' "3.78" ^ ^ xsd:double

Negative client property assertions

Negative data property assertions

FIG 7

DL query:	
Query (class expression)	
{P1}	
Execute	Add to ontology
Query results	
Super classes (7 of 7)	
☐ Clinical finding (finding) ☐ Disease (disorder) ☐ Patient (person) ☐ Person (person) ☐ SNOMED CT Concept (SNOMED RT + CTV3) ☒ Severe stenosis of aortic valve (disorder) ☐ owl:Thing	
Direct superclasses (2 of 2)	
☐ Patient (person) ☐ Severe stenosis of aortic valve (disorder)	

FIG 8




FIG 9




FIG 10





FIG 11



SYSTEM AND METHODS FOR KNOWLEDGE REPRESENTATION AND REASONING IN CLINICAL PROCEDURES

RELATED APPLICATION

[0001] This application claims the benefit of 21146514.4, filed Mar. 29, 2021, which is hereby incorporated by reference in its entirety.

FIELD

[0002] The present embodiments relate to an apparatus, a product, a method and a system for representation of medical knowledge for digital data processing.

BACKGROUND

[0003] Medical scientific research is evolving fast. Therefore, medical knowledge needs to be updated regularly to keep track with new developments.

[0004] In this regard, implementations of clinical guidelines exist, which may be used in clinical decision systems. Usually, these implementations are rule-based as, for example, by using decision tree implementations.

[0005] Generally, a decision tree can be linearized into a set of decision rules, where the outcome is the content of the leaf node, and the conditions along the path form a conjunction in the if-clause. In general, the rules have the form: “if condition1 and condition2 and condition3 then outcome”.

[0006] Decision rules can be generated by constructing association rules with the target variable on the right. They can also denote temporal or causal relations.

[0007] However, these implementations of clinical guidelines have limitations, because typically the logical model and knowledge representation of different implementations “below” the application layer is inhomogeneous and may differ from system to system. Therefore, a global collaboration system cannot be achieved.

[0008] In state-of-the-art, it is also known to use knowledge bases, using graph structured data models, like knowledge graphs. Knowledge graphs provide an alternative methodology to model rules for clinical decision support systems and there is much research that shows their advantage in resolving problems in the knowledge representation and automatic reasoning and decision support systems. Knowledge graphs may make use of intelligent machine learning algorithms and/or neural network technologies. Knowledge graphs may provide interfaces for associated software modules and systems. In particular, in the healthcare domain, a standardized ontology exists, the SNOMED ontology. However, for being able to use this technology of knowledge graphs in healthcare, generally, it is required that knowledge representation systems and related automatic support systems are internationally valid and may be used in different countries for different local requirements and with different local settings, including medical questions and/or specializations.

[0009] In the medical domain, often medical data (e.g., health-related measurements, vital data, anamnestic data) is missing but actually is necessary for providing a clinical decision. Further, not seldom, decisions have to be taken on an urgent basis for initiating life-threatening measures. Also, in this respect it would be helpful to provide an understanding of the reasoning and internal decisions (for example

which branches of a graph have been activated or which path in the graph has been taken), which have been taken by an automatic decision-support system, when providing a final result. In other words, it would be helpful to not only provide the final decision, but also to provide interim results to get more transparency for the interference machine and into the decision logic.

[0010] Generally, decision tree-based implementations have drawbacks. First, it is not possible to reach a conclusion (typically represented in traversing the tree down to a leaf) if it is not possible to provide all required variables or data.

[0011] Second, it is difficult or even impossible to provide other information than modeled within the decision tree, once it has been generated.

[0012] Third, as decision trees are “hard-coded,” it is difficult to update clinical knowledge and to provide machine learning capabilities.

SUMMARY AND DETAILED DESCRIPTION

[0013] Therefore, an object is to provide improved mechanisms and tools for knowledge representation for medical knowledge, which might be used by automatic clinical decision support systems. Further, automatic machine interference should be made as transparent as possible, and it should be possible to continuously update the knowledge base.

[0014] Further, it would be helpful to use the interference capabilities of knowledge graphs and also the inherent semantic relations of a certain ontology and not only to provide a clinical decision support system which points to an entry in the ontology data representation (e.g., SNOMED) as a result.

[0015] The above-mentioned object is solved with a method, apparatus, a system, and the non-transitory computer readable medium according to the attached independent claims. Further embodiments, features, and/or advantages are mentioned in the dependent claims.

[0016] In this respect, it has to be mentioned that the following description relates to an embodiment, represented in the method claim and its possible alternative embodiments. Features, advantageous embodiments, and optional alternatives, which are claimed and/or described with reference to the method may also be used in and thus transferred to the other claim types, like for example to the apparatus or to the computer program product and vice versa. From a computer scientist's point of view, a hardware implementation is to be regarded as equivalent to a software implementation. Therefore, any features mentioned with regard to the method (for example “storing”) may accordingly be transferred and applied to the hardware solution and will be implemented as hardware module with the respective functionality (“storage” configured with a storing functionality). For avoiding redundancies, any advantageous embodiments or features which are mentioned with respect to the method are not reiterated again for the apparatus or product claims.

[0017] Further, it is to be pointed out that software typically is modular in nature. Thus, a specific implemented feature which is mentioned in combination with a certain embodiment may also be combined with other features, even when mentioned in other embodiments. Accordingly, any feature may be combined with at least one other feature, which is claimed and/or described in this application.

[0018] With these embodiments, it is possible to implement procedural data, for instance clinical knowledge, e.g., in the form of guidelines onto existing ontologies so that it is possible to generate a new and updated ontology.

[0019] In particular, each guideline is treated as a set of edges between existing nodes, representing a new pathway in a knowledge graph. SNOMED ontology is taken as the base ontology due to it being the most widely used and most complete clinical terminology.

[0020] For processing these ontologies, standard ontology languages, like a Web Ontology Language (OWL), may be used, which is a family of knowledge representation languages for authoring ontologies. Further, a Resource Description Framework (RDF) may be used.

[0021] A knowledge base may be structured as a knowledge graph. The knowledge graph may be subject to a decision-tree inference system. In machine learning and decision analysis, a decision tree can be used to visually and/or explicitly represent decisions and decision making. It may use a tree-like model of decisions. A decision tree is, thus, a tool for deriving a strategy to reach a particular decision. Decision Trees are a supervised learning method used for classification or regression tasks. The goal is to create a model that predicts e.g., a health-related decision by learning simple decision rules inferred from the data features and procedural data. A tree can be seen as a piecewise constant approximation.

[0022] The embodiment is based on SNOMED. SNOMED is a multilingual thesaurus with an ontological foundation. The use of SNOMED makes clinical information available in a computable form and thus processable. Thus, clinical or healthcare data can be queried and used to trigger decision support rules and prompts. The hierarchies of SNOMED enable complex reasoning to support decision support rules. For example, in SNOMED the concept `!stroke!` is synonymous with `!cerebrovascular accident!` and subsumes all lower-level concepts including `!paralytic stroke!`, `!thrombotic stroke!` etc. This means that decision support queries are easier to develop and implement because they do not need to identify all the individual terms and codes which may be relevant. Moreover, SNOMED provides tools to extract their ontology to standard OWL language which can be interpreted by any semantic web reasoner.

[0023] SNOMED also has the concept of Expressions wherein pathways within their ontology can be represented in the form of equations and their own compositional grammar structure. For more details it is referred to SNOMED CT Compositional Grammar (2020), retrievable from <https://confluence.ihtsdotools.org/display/SLPG/SNOMED%2BCT%2BCompositional%2BGrammar>.

[0024] In an embodiment, procedural data, for instance, in the form of a guideline may first be mapped—preferably manually—to a set of SNOMED expressions. Each leaf in the guideline has one SNOMED expression that equates it to the input variables in the guideline. This allows its mapping to SNOMED concepts and attributes. No new classes or attributes are developed for this. The SNOMED expressions are then converted to OWL statements (statements in the Web Ontology Language) that are appended to the released SNOMED OWL file. This creates a new upgraded SNOMED+ OWL file which can form the backbone for a clinical decision support system.

[0025] In a deployment scenario, patient data in an electronic healthcare record (EHR) is mapped to a set of SNOMED concepts. There are existing tools that already do this like SemRep [see, for more details: Kilicoglu, H., Rosembat, G., Fiszman, M. et al. Broad-coverage biomedical relation extraction with SemRep. *BMC Bioinformatics* 21, 188 (2020). Retrievable from: <https://doi.org/10.1186/s12859-020-3517-7>]. Many hospitals also encode their data with SNOMED for standardization. An ontology reasoner can then use the guideline-enhanced SNOMED+ ontology provided by the tools of this application and infer if these patients belong to any specific classes. Ontology reasoners like e.g., Hermit [Hermit Reasoner (2013)] and compilers like Protégé [Protege OWL Environment (2011)] are well-established tools that are standard and can be used off-the-shelf. Descriptor Logic (DL Queries) can be used to also know direct and inferred super-classes that the patient belongs to.

[0026] A first aspect relates to a computer-implemented method for upgrading a medical knowledge base in a digital, clinical system. The method includes the acts of: providing a storage with a knowledge base, being a SNOMED knowledge base, in web ontology format (e.g., OWL or RDF or XML statements); receiving procedural data, representing clinical procedures for evaluation of a patient's health state (e.g., guideline data, in particular cardiological guideline data, like ECS data, ECS: European Society of Cardiology and/or clinical knowledge); mapping the received procedural data in a set of SNOMED expressions; converting the SNOMED expressions into statements in the web ontology format; and upgrading the SNOMED knowledge base with the received procedural data by adding the statements in the SNOMED knowledge base for providing a processable file with an upgraded version of the SNOMED knowledge base.

[0027] In a preferred embodiment, the knowledge base includes or is based on a medical semantic network and in particular may apply a graph-based decision tree. For more technical details with respect to graph-based decision trees and its functionality it is referred to Corbett D. R. (2008) *Graph-Based Representation and Reasoning for Ontologies*. In: Fulcher J., Jain L. C. (eds) *Computational Intelligence: A Compendium*. Studies in Computational Intelligence, Vol. 115. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-540-78293-3_8.

[0028] The invention provides significant advantages.

[0029] First, it is possible to provide a result, for example, a decision with respect to a health state of a patient, even in the case data is missing or is not complete. This is possible because the embodiment is based on a graph-based decision tree. Due to the acyclic and non-directional nature of graphs as opposed to (normal) decision trees that can only move forward if all conditions of previous acts have been met, the graph-based decision tree according to the solution presented herein may also provide data in difficult scenarios, with e.g., missing or incomplete data.

[0030] Second, the inference for the decision may be made transparent. In particular, upon user request, he or she is informed about automatic decision processing, i.e., which branches and paths in a graph-based decision tree have been activated. Since a graph is based on parent-child-attribute relationships, the end leaf node classification can be traced back to the uppermost parents in the graph of all variables that led to it. In this way, the trace provides a method for the user to know the path taken by the graph to arrive at the

classification. It is more explainable. For more details, it is referred to the Figures, in particular to the description of FIG. 2.

[0031] Third, it is very easy to upgrade the knowledge base and/or the rules for automatic decision execution. Since a graph is acyclical and non-directional, if a relation between two variables is established, the link can just be added. The entire decision tree doesn't have to be modified and tested. The ontology reasoner algorithm will function as usual and take into consideration the new linked entity in its reasoning. In this way, it is very easy to add/delete/modify relations without requiring a redesign of the entire decision tree. For example, if "hyperthyroidism" has to be added as a requirement for "moderate AS", the SNOMED concept for positive hyperthyroidism has to be just added as attribute to moderate AS.

[0032] With this embodiment, specific types of inference are possible and processable, which would not be possible with a normal decision tree. For example, the SNOMED concept for Aortic Stenosis may be linked to other SNOMED concepts regarding LVEF, AVA, Vmax and Delta PM (Pressure Gradient) through SNOMED relationships, wherein LVEF—Left Ventricle Ejection Fraction, AVA (Aortic Valve Area), Vmax—Maximal Flow Velocity in Valve, Delta PM (Pressure Gradient). In SNOMED ontology, since the SNOMED concept for Aortic Stenosis is defined by LVEF values as well, a SNOMED-based graph can detect automatically that the values are discordant and that the patient could be having "pseudosevere AS". Thus, clinical semantic knowledge embedded in SNOMED is used for more accurate classifications in this technology.

[0033] In another preferred embodiment, the knowledge base may be extended by further ontologies. For instance, ICD codes used for classification of diseases may also be used and applied in the form of an ontology, i.e., more detailed versions of a particular disease are children of the parent abstract disease. ICD codes are used in electronic health record/EHR systems to code the disease that a patient had and also for billing purposes. The ICD ontology can be linked to the SNOMED ontology in the same manner as concatenating two graphs. In such a scenario, the reasoning can be done not just based on SNOMED input variables but also on ICD input variables, thus enriching the pathway. Classification output can also be either SNOMED or ICD.

[0034] In another preferred embodiment, a specific patient instance may be applied to the upgraded processable file for inference of the patient's health state by loading the upgraded processable file in an ontology reader. In particular, the ontology reader may be configured to implement a reasoning algorithm. In particular, the reasoning algorithm may be or may include of classification algorithm in order to classify a patient's health state into disease categories and/or measure categories, representing medical measures to be performed.

[0035] In another preferred embodiment, the processable file is capable to provide result data even if a parameterization is incomplete, and/or parameters are inconsistent. For example, in patients with cardiac catheterization providing direct measurement of the pressure gradient across the valve or by alternative methodologies, like patient-based physiological modelling, measurement of blood velocity through Doppler signal processing may not be required to distinguish between low and high gradient AS. In a decision tree method, if the blood velocity Vmax is not present, it

wouldn't function to provide a result, as that is the first entry point. But in a graph, if the AVA value is directly provided, it can skip the first requirement and continue. This is, because there is no precondition or direction in graph-based reasoning that blocks classification. If values can be provided within the pathway, the graph can continue from that point.

[0036] In still another preferred embodiment, the method may further include the act of: generating a virtual representation of the decision tree processing. The virtual representation may be stored and provides an automatically generated documentation of the inference and decision tree execution, so that the processing becomes more transparent and interim results may be documented as well. Alternatively, or in addition, the virtual representation may be used to automatically validate and, if needed, correct the processing. Moreover, by using the virtual representation, a secondary system may automatically validate and/or infer the processing to be carried out.

[0037] In another preferred embodiment, the knowledge base is augmented by numerical measurement values acquired by a set of medical devices. The medical devices may include a temperature sensor, a device for measuring blood composition and values and/or acquisition devices, in particular devices for acquiring physiological parameters or others.

[0038] In still another preferred embodiment, the knowledge base is augmented by medical image data acquired by a set of medical imaging devices, like CT, MRI, ultrasound, and/or others.

[0039] In still another preferred embodiment, the upgraded SNOMED knowledge base is used in a clinical decision support system (CDS) for providing a classification result dataset. The classification task is determined on a case-by-case basis and according to the respective setting. The classification may e.g., be in two or more different classes, like healthy and non-healthy.

[0040] In still another preferred embodiment, the classification result dataset includes a prediction for a health-related measure and/or process and/or a patient's health state, a confidence range, further clinical measures and/or an inference dataset, representing a trace of an automated reasoning leading to the classification result dataset. The technical effect of this features is to enhance transparency of automatic inference.

[0041] Preferably, the method may include a machine learning algorithm. The machine learning algorithm may be configured to include instructions for processing input data to obtain output data and at least some of these instructions are set by using a set of training data and a training algorithm. The machine learning algorithm may be a supervised learning algorithm. It may be based on a convolutional neural network and/or a deep neural network, a decision tree, a random forest, a support vector machine (SVM), or other architectures.

[0042] Cumulatively or in addition, the machine learning algorithm may respectively use or include a decision tree that processes input data to "decide", e.g., provide an estimate, of a patient's health state, in particular providing an estimate of which kind of device-based measure, procedure or care (e.g., with an imaging device) is necessary for the patient.

[0043] In still another preferred embodiment, a trained graph-based convolutional neural network may be used for

performing a data processing task, like a classification task with respect to the patient's health status. Alternatively, or in addition, a deep neural network (DNN) and in particular a trained graph-based convolutional neural network (CNN) may be used for estimating missing or incorrect data.

[0044] The training data of the neural network may be an instance of an incomplete data set (very different medical and/or healthcare data types and/or parameters are missing) and an instance of a correct and complete dataset, which is related and associated to the incomplete data set and/or manually checked. After training, the neural network is then used to automatically complement or complete insufficient and/or inconsistent and/or partial data sets.

[0045] In still another preferred embodiment, self-explanation techniques are applied for explaining convolutional neural network interference. Preferably, the self-explanation techniques may be based on a layer-wise relevance propagation technique, LRP).

[0046] Interpretability is especially important in applications such as medicine, where the reliance of the model on the correct features must be guaranteed. Therefore, transparency is important and has the technical effect that it is possible to provide an explanation of algorithmic (and thus automatic and "hidden") decisions. With this feature, it is possible to explain why a model is making the respective predictions. This serves to improve reliability.

[0047] An approach to explaining the prediction of the deep neural network (DNN) is to make explicit use of its graph structure and proceed as follows: It is started at the output of the network. Then, it is moved in the graph in reverse direction, progressively mapping the prediction onto the lower layers. The procedure stops once the input of the network is reached. Layer-wise mappings can be engineered for specific properties.

[0048] Layer-wise relevance propagation (LRP), for example, is applicable to general network structures including DNNs and kernels. The layer-wise mappings are designed to ensure a relevance conservation property, where the share of received by each neuron is redistributed in same amount on its predecessors. The injection of negative relevance may be controlled by hyperparameters. For more details of the Layer-wise relevance propagation (LRP) it is referred to S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, W. Samek, "On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation", PLoS ONE, 10 (7) (2015), Article e0130140.

[0049] In a preferred embodiment, a reasoner is used. The provided processable file may be forwarded and provided to the reasoner. The reasoner is configured to be sound, complete, and terminating: i.e., all entailments it finds do indeed hold, it finds all entailments that hold, and it always terminates. An "inference engine" may rely on a reasoning algorithm, implemented in a reasoner. The reasoner may be used for different purposes, like classification or other reasoning services. The invention makes use of automated reasoners such as Pellet, FaCT++, HerMIT, ELK etc. to take a collection of axioms written in OWL and offer a set of operations on the ontology's axioms.

[0050] In still another preferred embodiment, the method includes accessing an extraction tool for extracting selected features of a SNOMED ontology.

[0051] The order, according to which the acts of the method of the present invention are described in the present

specification, does not necessarily reflect the chronological order, according to which said acts are carried out.

[0052] According to another aspect, an apparatus is provided for upgrading a medical knowledge base in a digital, clinical system. The apparatus includes: a first input interface to a storage for storing a knowledge base, being a SNOMED knowledge base, in web ontology format; a second input interface, configured to receive procedural data, representing clinical procedures for evaluation of a patient's health state; a processing unit (processor) for mapping the received procedural data in a set of SNOMED expressions; wherein the processing unit is further adapted for converting the SNOMED expressions into statements in the web ontology format; and wherein the processing unit is configured for upgrading the SNOMED knowledge base with the received procedural data by adding the statements in the SNOMED knowledge base for providing a processable file with an upgraded version of the SNOMED knowledge base.

[0053] In another aspect, a reasoning system is provided. The reasoning system includes: an apparatus, as mentioned above and the reasoner (inference engine). The reasoning system may be or may be integrated in a clinical decision support system (CDS).

[0054] According to another aspect, a computer program product includes program elements that induce a computer to carry out the acts of the method for upgrading a medical knowledge base in a digital, clinical system according to any of the preceding method claims, when the program elements are loaded into a memory of the computer.

[0055] According to another aspect, a non-transitory computer-readable medium is provided. The non-transitory computer readable medium stores program elements that can be read and executed by a computer in order to perform acts of the method for upgrading a medical knowledge base in a digital, clinical system according to any of the preceding method claims, when the program elements are executed by the computer.

[0056] The realization by a computer program product and/or a computer-readable medium has the advantage that already existing servers and/or clients can be easily adopted by software updates in order to work as proposed by the invention.

[0057] In the following, a definition of terms used within this application is given.

[0058] A knowledge base is to be construed as a storage with a knowledge representation stored therein. A knowledge base may be a knowledge graph. A knowledge base serves as storage technology used to store and to make retrievable complex structured and unstructured information used by a computer system. An object model or ontology may serve as representation of knowledge. A knowledge base may be construed as a representation of heuristic and factual information, often in the form of facts, assertions, and deduction rules. A knowledge base is usually used together with an inference engine (also called reasoner), which is a mechanism, playing the role of an interpreter, that applies the knowledge as represented in a suitable way to achieve results (for instance, automatic decisions, classification results etc.). Usually, a knowledge-based system further includes interfaces for data exchange, in particular a user interface, with a mechanism that transfers queries from

and answers to the user, sometimes seeking additional information for the Inference Engine. This includes explanation facilities for the user.

[0059] In particular, the knowledge base is a SNOMED knowledge base and is provided in a web ontology format. In a knowledge base, the knowledge representation should be a hybrid of many individual representations, such as frames, database facts, and deduction rules. Often, knowledge units expressed in one of these languages must be aggregated to the hybrid knowledge representation of one knowledge source which in turn must be further aggregated to a globally consistent knowledge base.

[0060] Within this application, the term “knowledge graph” is to be construed as that it acquires and integrates information into an ontology and applies a reasoner to derive new knowledge. For further definition it is referred to “Towards a Definition of Knowledge Graphs”, L. Ehrlinger, W. Wöß, SEMANTICS 2016: Posters and Demos Track Sep. 13-14, 2016, Leipzig, German. Thus, in this application a knowledge graph is somehow superior and more complex than a “normal” knowledge base (e.g., an ontology) because it is structured and formalized that it can be used by a reasoning engine (reasoner) to generate new knowledge and integrates one or more information sources.

[0061] SNOMED (or SNOMED CT, CT for clinical terms) is a systematically organized computer processable collection of medical terms providing codes, terms, synonyms and definitions used in clinical documentation and reporting. SNOMED is considered to be the most comprehensive, multilingual clinical healthcare terminology in the world. The primary purpose of SNOMED is to encode the meanings that are used in health information and to support the effective clinical recording of data with the aim of improving patient care. SNOMED provides the core general terminology for electronic health records. SNOMED comprehensive coverage includes: clinical findings, symptoms, diagnoses, procedures, body structures, organisms and other etiologies, substances, pharmaceuticals, devices and specimens. In this application SNOMED serves as a reference terminology in a multilingual form (i.e., enabling multilingual use).

[0062] Generally, a storage may include volatile primary memory (e.g., a RAM, a DRAM, a SRAM, a CPU cache memory or the like) and/or non-volatile primary memory (e.g., a ROM, a PROM, an EPROM or the like). In particular, the volatile primary memory may consist of a RAM. For instance, the volatile primary memory temporarily holds program files for execution by the processing unit and related data and the non-volatile primary memory may contain bootstrap code for the operating system of the data processing system. The storage may further include a further memory, which may store the operating system and/or the instructions of the algorithms used to carry out the method of the present invention, in particular upgrading the SNOMED knowledge base. Moreover, the further memory may store a computer program product including instructions which, when the computer program product is executed by the processing unit, cause the processing unit to carry out the method according to the present invention.

[0063] Upgrading the knowledge base means to update the data storage with new data. This may preferably be done on a regular basis. This includes to store additional data into the knowledge base in a related form, which means that new

data is assigned and associated to respective corresponding entries in the knowledge base. So, new data is processable.

[0064] Procedural data is a set of digital data. Procedural data may be stored in a distributed manner at different storage locations. Procedural data represent clinical procedures for evaluation of a patient’s health state or for measures to be taken in this respect, e.g., initiating medical examinations or evaluations (like image acquisition procedures etc.).

[0065] SNOMED expressions are portions in digital form which are processable in SNOMED. SNOMED expressions may be understood as a set of relationships between SNOMED Concept Codes. This basically is identical to an equation with SNOMED Concept codes as variables.

[0066] Statements in the web ontology format or language, in particular OWL statements, are statements to be processed digitally. The Web Ontology Language (OWL) is designed for use by applications that need to automatically process the content of information by algorithms and machines instead of just presenting information to humans. OWL facilitates greater machine interpretability of Web content than that supported by XML, RDF, and RDF Schema (RDF-S) by providing additional vocabulary along with a formal semantics. OWL has three increasingly expressive sublanguages: OWL Lite, OWL DL, and OWL Full, which all can be used in this invention. OWL can be used to explicitly represent the meaning of terms in vocabularies and the relationships between those terms. This representation of terms and their interrelationships is called an ontology.

[0067] Procedural data represent medical or clinical procedures or actions, like e.g., measurements with medical devices. Procedural data may include guideline data.

[0068] Wherever not already described explicitly, individual embodiments, or their individual aspects and features, described herein can be combined or exchanged with one another without limiting or widening the scope of the described invention, whenever such a combination or exchange is meaningful and in the sense of this invention.

[0069] The order, according to which the acts of the method are described in the present specification, does not necessarily reflect the chronological order, according to which said acts are carried out.

[0070] Other variations to the disclosed embodiments can be understood and effected by those skilled in the art in practicing the claimed invention, from a study of the drawings, the disclosure, and the appended claims. In the claims, the word “comprising” does not exclude other elements or acts, and the indefinite article “a” or “an” does not exclude a plurality.

[0071] Advantages which are described with respect to a particular embodiment of present invention or with respect to a particular figure are, wherever applicable, also advantages of other embodiments of the present invention.

[0072] It shall be understood that a preferred embodiment of the present invention can also be any combination of the dependent claims or above embodiments or features with the respective independent claim(s).

[0073] These and other aspects of the invention will be apparent from and elucidated with reference to the embodiments described hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0074] FIG. 1 shows an embodiment of the architecture for implementing the method for upgrading a SNOMED knowledge base;

[0075] FIG. 2 is an example for a stepwise integrated approach for the assessment of Aortic stenosis severity (left) and focus view (right side);

[0076] FIG. 3 is an example for a simplified guideline as procedural data;

[0077] FIG. 4 shows another example of trimmed SNOMED ontology;

[0078] FIG. 5 shows an example with three patient instances loaded onto a reasoner, working upon the upgraded knowledge base, being provided according to the invention;

[0079] FIG. 6 shows an example result, which has been classified by the reasoner to both moderate stenosis and the adult classes;

[0080] FIG. 7 depicts examples for direct and indirect SNOMED super-classes of a patient using descriptor logic query;

[0081] FIG. 8 is a flow chart of the method according to a preferred embodiment;

[0082] FIG. 9 is another flow chart of optional procedures which may be applied on top of the general procedure represented in FIG. 8;

[0083] FIG. 10 is another flow chart of optional procedures that may be applied on top of the general procedure represented in FIG. 8;

[0084] FIG. 11 is a block diagram of an apparatus for executing the method according to a preferred embodiment.

DETAILED DESCRIPTION OF THE DRAWINGS

[0085] The present embodiments relate to a method and apparatus and related systems for upgrading a medical knowledge base with procedural data, like guideline data.

[0086] As can be seen in FIG. 1, procedural data PD should be applied to a knowledge base, being a SNOMED knowledge base. The content of the knowledge base is converted to OWL statements, to provide an upgraded or enhanced SNOMED ontology OWL file, which may be subject to further digital and automatic data processing. It is possible to add classes and relations and/or other entities according to the standard OWL language. Based upon the enhanced knowledge base, a specific patient instance may be added using an interface, for example an OWL API for inferring result data, by using an ontology reasoner R (see FIG. 11). The result may be interpreted by a clinical decision system CDS later on.

[0087] Without loss of generality, for instance, the translation of the clinical guidelines for the assessment of aortic valve stenosis into the upgraded knowledge base which has been provided with this embodiment and which may be denoted as SNOMED+ ontology.

[0088] An example is depicted in FIG. 2 as an example of a graphical representation of the recommended clinical decision process. As can be seen, the sequence of acts may be identified in the procedural data PD. For example, if the classification is "Moderate AS" (AS-Aortic Stenosis), the trace can read as follows: "Moderate AS (Child)->AS(parent)->[AVA(attribute), Flow velocity(attribute), Pressure Gradient(attribute)]". This means that it was Moderate AS because of variables AVA, Flow Velocity and Pressure

Gradient that affected its parent AS. The trace will become more useful for much deeper graphs.

[0089] For example, the SNOMED concept for Aortic Stenosis is linked to other SNOMED concepts regarding LVEF, AVA, Vmax and Delta PM (Pressure Gradient) through SNOMED relationships, wherein LVEF—Left Ventricle Ejection Fraction, AVA (Aortic Valve Area), Vmax—Maximal Flow Velocity in Valve, Delta PM (Pressure Gradient). In the following scenario: Vmax and Delta PM are at thresholds and discordant or inconsistent to each other i.e., Vmax is <4 and Pm>40. In such a case a decision tree doesn't "know" which way to proceed. Even if the decision tree is programmed to take the left subtree, the AVA measurement could also be at threshold. If it is >1, the decision tree will say "Moderate AS". Now, LVEF is a measurement independently taken from the image and is not derived from Vmax, AVA, or Pm. Hence, the error is not propagated. If the LVEF value is much less than 50, a clinician will know that the patient could have "pseudosevere AS" and not "moderate AS." This could not be reflected in a usual decision tree as the LVEF value only comes into play down the tree and cannot be used to correct starting variables like AVA, Vmax, or Delta PM. In SNOMED ontology, since the SNOMED concept for Aortic Stenosis is defined by LVEF values as well, a SNOMED-based graph can detect automatically that the values are discordant and that the patient could be having "pseudosevere AS." Thus, clinical semantic knowledge embedded in SNOMED is used for more accurate classifications in this technology.

[0090] The recommended decision process for the classification of the disease (here: aortic stenosis, AS) is based on the evaluation of multiple clinical statuses, which are combined into a final diagnosis. Each clinical status may be evaluated based on direct or indirect clinical measurements (e.g., measurement of maximal velocity Vmax through the aortic valve). SNOMED expressions can describe each path in the decision process through SNOMED concepts and attribute.

[0091] One advantage of this approach is that it allows to leverage equivalency between concepts through nested SNOMED expressions. For instance, 'high flow status' can itself be defined as an expression, to be evaluated based on a set of input variables (concepts and attributes). For the specific case of high flow status, conditions associated to it include anaemia, hyperthyroidism, presence of arterio-venous shunts. Assessment of high flow status can then be achieved in one or multiple ways, based on the same or different clinical guidelines, so that in the deployment scenario the decision pathway is automatically cleared based on available clinical information.

[0092] In the same spirit, a graph-based decision tree can also be augmented with decision pathways relying on additional information leading to the same diagnosis. For instance, in patients with cardiac catheterization providing direct measurement of the pressure gradient across the valve, measurement of blood velocity through Doppler signal processing may not be required to distinguish between low and high gradient AS (AS: aortic stenosis). Similarly, alternative decision processes can be designed based on alternative methodologies such as non-invasive estimation of blood pressure gradients based on patient-tailored physiological modeling. Such approaches may themselves be represented by SNOMED expressions and rely on additional concepts and attributes (for instance, blood pressure can be

estimated by physiological modeling based on 3D imaging and cuff blood pressure measurements). In this context, physiological modeling refers to estimation of physiological signals or phenomena based on clinical data, using various approaches including computational modeling or machine trained models.

[0093] The reasoning system based on an upgraded SNOMED+ ontology described can be based on existing guidelines. Procedural data PD, for example, guideline data, can be extracted for instance from published clinical recommendations automatically. For this purpose, an extraction tool ET may be applied as a “crawler” for searching configured databases on a regular basis. Procedural data PD may also be based on clinical practice as defined within specific communities or healthcare systems (best practice, represented in a digital format). It can also be used to expand existing decision processes, for instance by defining specialized decision processes for subgroup of patients. One possible implementation of specialized decision processes is based on explicitly conditioning different SNOMED expressions based on patient characteristics or comorbidities. A simple example is adapting the thresholds in the decision process depicted in the picture above based on patient anatomical conditions (patients with small LVOT diameters, or generally patients of small stature are expected to have small valve area in presence of low-pressure gradient in normal conditions). Another implementation is based on explicitly enriching the decision process by including non-imaging clinical observations (patient info, symptoms, clinical report of examination) and designing additional or alternative decision pathways based on those observations. One example of this approach is a decision process looking at patients for which diagnosis of aortic stenosis is difficult purely based on the standard measurements recommended by the guidelines—such as patients with stenotic valve in presence of low-pressure gradient, small aortic valve area and preserved ejection fraction. In this subgroup, collection of symptoms, personal data and additional measurements (e.g., calcium score) are all useful for the final determination.

[0094] One important application of the proposed system and method is that the system may serve as the backbone of interpretable decision support systems CDS.

[0095] The user of the system receives an estimated disease condition for a given patient, together with a trace of the reasoning leading to the decision. This can be augmented with confidence intervals which can be associated to each SNOMED expression, for instance as a function of the number of available input variable or the degree of certainty on each measurement.

[0096] The proposed system models the relationship between SNOMED concepts/attributes and the diagnosis through pre-defined SNOMED expressions. It can be expanded to create new expressions based on available concepts.

[0097] In a preferred embodiment, graph convolutional neural networks can be applied to model the relationship between SNOMED concepts and given disease labels, as following. An upgraded SNOMED+ ontology as being provided by the solution, presented herein, as processable file PF, can be represented as a graph, with SNOMED expressions encoding guidelines/decision processes being represented by sub-graphs. A training database is defined as a multitude of graphs, each corresponding to a subject (after

mapping subject data to SNOMED+); each subject is associated to a disease label. Without loss of generality, for the example of aortic stenosis, the label can be: no disease, moderate disease, severe disease. Graph convolutional neural network process input data with a graph structure by computing predictive feature in each graph node based on information from the node neighbors. This allows such networks to be robust to changing graph architectures (corresponding to missing information) and to be sensitive to data patterns associated to nodes that are explicitly linked together. They are therefore suited for use on graphs defined based on ontologies such as SNOMED, in which node connections are based on semantic information based on medical knowledge. Finally, the neural network can be trained to perform a classification task by a classification algorithm CA, each class representing a disease condition. **[0098]** Importantly, techniques to explain the behavior of the neural network (such as layer-wise relevance propagation) can be used to discover what SNOMED concepts or combination thereof are relevant for correct disease diagnosis in the given training population.

Demo Use Case:

[0099] For demo purposes, the standard SNOMED ontology is augmented with a modified version of the above ESC guideline (ESC: European Society of Cardiology) along with error margin incorporation.

[0100] In the following, acts which are undertaken are listed:

[0101] 1) The guideline is simplified even further for the purpose of demo and is shown in a graphical simplified schematic representation of FIG. 3.

[0102] 2) The SNOMED ontology is also trimmed to only include concepts and relations that are relevant for this demo (FIG. 4).

[0103] 3) The SNOMED ontology file is then converted to standard OWL format using the SNOMED OWL Toolkit, for example, SNOMED OWL Toolkit (2020). Retrieved from <https://github.com/IHTSDO/snomed-owl-toolkit>.

[0104] 4) The simplified guideline is converted to the following SNOMED expressions so as to classify the three leaf nodes (Severe stenosis, Moderate Stenosis and No Stenosis or Healthy Adult). SNOMED compositional grammar is used to create the expressions and its description is beyond the scope of this document.

[0105] 836482000|Severe stenosis of aortic valve (disorder)|==(252065006|Peak arterial velocity (observable entity)|:732944001|Has presentation strength numerator value (attribute)|>= #4+251081004|Cardiovascular pressure gradient (observable entity)|:732944001|Has presentation strength numerator value (attribute)|>= #40

[0106] 836481007|Moderate stenosis of aortic valve (disorder)|==(252065006|Peak arterial velocity (observable entity)|:732944001|Has presentation strength numerator value (attribute)|< #4+251081004|Cardiovascular pressure gradient (observable entity)|:732944001|Has presentation strength numerator value (attribute)|< #40+251011009|Aortic valve area (observable entity)|:732944001|Has presentation strength numerator value (attribute)|> #1

[0107] 102512003|Well adult (finding)|==836481007|Moderate stenosis of aortic valve (disorder)|==(252065006|Peak arterial velocity (observable entity)|:732944001|Has presentation strength numerator value (attribute)|< #4+251081004|Cardiovascular pressure gradi-

ent (observable entity):732944001|Has presentation strength numerator value (attribute):<#40+251011009|Aortic valve area (observable entity):732944001|Has presentation strength numerator value (attribute):<= #1

[0108] 5) The SNOMED expressions are then converted to a set of OWL statements and added to the trimmed SNOMED OWL file. Expressions can be written as equivalent classes in OWL RDF/XML syntax.

[0109] 6) The input numerical variables are redefined to have maximum and minimum range values to account for the assumed 5% error rate.

[0110] 7) 3 Patient instances are then added to the same OWL file to have values that lead to the three different classifications.

[0111] 8) The processable file PF with the modified SNOMED+ OWL file can then be loaded onto any ontology reader R and then any standard reasoner can be applied. Protégé may be used for this purpose and the inbuilt Hermit reasoner may be used, too.

[0112] As seen in FIG. 5, the reasoner R automatically classifies the patient instances to one of the three stenosis classes.

[0113] 9) For patient P2 (FIG. 6), the reasoner R (not shown in FIG. 6, however see FIGS. 1 and 11) infers both Moderate Stenosis and Well Adult SNOMED classes since the Aortic Valve Area (AVA) was at the threshold of 1 cm². Considering the error rate of 5%, P2 could be classified as having either of these conditions. Such a result of opposing classes can prompt the physician to confirm the value of AVA (aortic valve area) before further decisions.

[0114] 10) One can also use the inherent SNOMED relations and standard ontology Description Logic queries to identify other classes to which the patient belongs (FIG. 7). In this case, the reasoner R infers automatically that since P1 has Severe Stenosis, he/she should have a disease and thus also a clinical finding as super-classes. This is because in SNOMED, Sever Stenosis concept is classified as a disease and a disease is classified as a clinical finding. This can be extended to add other linked conditions, diseases or treatment options in the SNOMED ontology.

[0115] SNOMED's semantic network can be leveraged to make the guidelines better related to the latest clinical knowledge. It also allows guidelines to be implemented without the need to re-define all possible concepts. SNOMED is not just used for reference literature, but its semantic network, relations and concepts are used in addition to the guideline or other procedural data in the same upgraded ontology.

[0116] In this embodiment, an existing ontology is used as the base for adding guidelines. Standard ontology languages are used to facilitate easy plug-and-play of SNOMED versions and interpretation by various frameworks like Python/Java/Protege.

[0117] Other terminologies, like ICD-10, CPT, LOINC, can be linked to the ontology too using existing released mappings.

[0118] SNOMED expressions are leveraged to map to existing concepts and provide a method for cross-validation by multiple clinicians.

[0119] A graph-based decision tree also means that not all input variables are required to make decisions and it doesn't get blocked at any level. Definitions of relations between variables, either via SNOMED ontology or user-defined, can

be used to assume values for missing variables. Confidence measures can be implemented to give probabilistic decisions depending on availability of data points.

[0120] As shown above, ontology-based reasoning can also show smarter decisions in threshold cases as opposed to deterministic outputs.

[0121] Changes in clinical knowledge can be updated either via new releases of SNOMED or adding them as new relations or concepts manually in the SNOMED OWL. A re-definition of the rules-engine is not needed, thus saving time and cost.

[0122] FIG. 8 shows a flowchart of a method for upgrading and/or operating a medical knowledge base. After start of the method, in act S1, a storage is provided with a knowledge base, being a SNOMED knowledge base, in a web ontology format (OWL). In act S2 procedural data, representing clinical procedures for evaluation of the patient's health state, like guideline data, are received. In act S3 the received procedural data is mapped to a set of SNOMED expressions. In act S4 the SNOMED expressions are converted into statements in the that ontology format. In act S5, the SNOMED knowledge base is upgraded with the received procedural data. This is preferably executed by adding the statements in the SNOMED knowledge base for providing processable file PF in act S6 with an upgraded version of the SNOMED knowledge base.

[0123] As can be seen in FIG. 9, further algorithms may be applied upon the upgraded version of the SNOMED knowledge base in the processable file PF. For example, after having provided the upgraded version of the SNOMED knowledge base in act S6, a classification algorithm CA may be applied in order to classify a patient's health state (for example in healthy or disease with different severities). After having provided the classification result, the method may end.

[0124] Another example of the application of a subsequent algorithm is shown in FIG. 10. After having provided the upgraded version of the SNOMED knowledge base in the processable file PF in act S6, a reasoning algorithm RA, implemented in the reasoner R, may be applied in order to show and make transparent the interference which has been used to come to the result, in particular to the classification result. The result is provided in a processable file, too. The processable file may be back propagated to the reasoner R.

[0125] FIG. 11 shows the architecture of an automatic reasoning system with an apparatus, which is configured to execute the method for upgrading a medical knowledge base. The apparatus includes or is in data exchange with a first memory depicted in FIG. 11 with reference MEM1, in which the knowledge base, in particular the SNOMED knowledge base, is stored. A second memory MEM2 is provided for storing guideline data or other procedural data PD. Depending on the specific embodiment, it is possible that the first and the second memory are implemented on the same or a different entities or units. The central processing unit CPU (processor) includes a first interface i1 for providing data exchange to the SNOMED knowledge base, in a web ontology format. A second input interface i2 is configured to receive procedural data PD, representing clinical procedures for evaluation of a patient's health state, e.g., guideline data. The processing unit CPU is configured for mapping the received procedural data in a set of SNOMED expressions. The processing unit CPU is further adapted for converting the SNOMED expressions into statements in the

web ontology format and for upgrading the SNOMED knowledge base with the received procedural data by adding the statements in the SNOMED knowledge base for providing a processable file PF with an upgraded version of the SNOMED knowledge base. This can be seen in FIG. 11, the apparatus is shown in the middle part and mainly includes the two interfaces i1, i2 and the CPU for data processing, in particular for executing the machine learning algorithms for providing the processable file PF.

[0126] The processable file PF is provided as output of the method and/or apparatus. The processable file PF may be processed by a reasoner R, which acts as inference engine on the processable file PF.

[0127] In a first embodiment, the reasoning system may include the apparatus and the reasoner R. The reasoner R may interact with an external clinical decision support system CDS, which may include a user interface UI for user interaction. This embodiment is shown in FIG. 11 with dotted lines.

[0128] In a second embodiment, the reasoning system itself is part of the clinical decision support system CDS. This is reflected in FIG. 11 with dashed lines.

[0129] As can be seen in FIG. 11, the processing unit CPU may include an extraction to ET for extracting selected features of a SNOMED ontology. This feature improves performance of the method and may focus the upgrading process upon only relevant data. With this feature, execution time for the upgrading process may be reduced.

[0130] Any reference signs in the claims should not be construed as limiting the scope.

[0131] Wherever not already described explicitly, individual embodiments, or their individual aspects and features, described in relation to the drawings can be combined or exchanged with one another without limiting or widening the scope of the described invention, whenever such a combination or exchange is meaningful and in the sense of this invention. Advantages which are described with respect to a particular embodiment of present invention or with respect to a particular figure are, wherever applicable, also advantages of other embodiments of the present invention.

1. A computer-implemented method for upgrading a medical knowledge base in a digital, clinical system, the method comprising:

- providing a storage with a knowledge base, being a SNOMED knowledge base, in a web ontology format;
- receiving procedural data representing clinical procedures for evaluation of a patient's health state;
- mapping the received procedural data in a set of SNOMED expressions;
- converting the SNOMED expressions into statements in the web ontology format;
- upgrading the SNOMED knowledge base with the received procedural data by adding the statements in the SNOMED knowledge base for providing a processable file with an upgraded version of the SNOMED knowledge base, wherein the knowledge base comprises a graph-based decision tree; and
- generating a virtual representation of the decision tree processing.

2. The method according to claim 1, wherein the knowledge base is extended by further ontologies.

3. The method according to claim 1, wherein the upgraded SNOMED knowledge base is used for classification of a patient's health state by loading the provided processable

file in an ontology reader, wherein the ontology reader is configured for applying a classification algorithm.

4. The method according to claim 1, wherein a specific patient instance is applied to the processable file for inference of the patient's health state by an ontology reader for applying a reasoning algorithm.

5. The method according to claim 1, wherein the processable file provides result data even when a parameterization is incomplete and/or parameters to be processed are inconsistent.

6. The method according to claim 1, wherein the knowledge base is augmented by numerical measurement values acquired by a set of medical devices.

7. The method according to claim 1, wherein the knowledge base is augmented by medical image data acquired by a set of medical imaging devices.

8. The method according to claim 1, wherein the upgraded SNOMED knowledge base is used in a clinical decision support system, providing a classification result dataset.

9. The method according to claim 8, wherein the classification result dataset comprises a prediction for a patient's health state, a confidence range, further clinical measures, and/or an inference dataset, representing a trace of an automated reasoning leading to the classification result dataset.

10. The method according to claim 9, wherein self-explanation techniques are applied for explaining convolutional neural network inference.

11. The method according to claim 1, further comprising accessing an extraction tool for extracting selected features of a SNOMED ontology.

12. An apparatus for upgrading a medical knowledge base in a digital, clinical system, the apparatus comprising:

- a first input interface configured to provide a knowledge base, being a SNOMED knowledge base, in web ontology format;
 - a second input interface configured to receive procedural data, representing clinical procedures for evaluation of a patient's health state;
 - a processor configured to map the received procedural data in a set of SNOMED expressions;
- wherein the processor is further configured to convert the SNOMED expressions into statements in the web ontology format;
- wherein the processor is configured to upgrade the SNOMED knowledge base with the received procedural data by adding the statements in the SNOMED knowledge base to provide a processable file with an upgraded version of the SNOMED knowledge base;
- wherein the knowledge base comprises a graph-based decision tree; and

wherein the processor is further configured to generate a virtual representation of the decision tree processing.

13. The apparatus according to claim 12, wherein the knowledge base is extended by further ontologies.

14. The apparatus according to claim 12, wherein the processor is configured to use the upgraded SNOMED knowledge base for classification of a patient's health state by loading the provided processable file in an ontology reader, wherein the ontology reader is configured to apply a classification algorithm.

15. The apparatus according to claim 12, wherein the processor is configured to apply a specific patient instance to

the processable file for inference of the patient's health state by an ontology reader for applying a reasoning algorithm.

16. The apparatus according to claim **12**, wherein the processable file is configured to provide result data even when a parameterization is incomplete and/or parameters to be processed are inconsistent.

17. A non-transitory computer-readable storage medium on which program elements are stored that can be read and executed by a computer to upgrade a medical knowledge base in a digital, clinical system, the program elements comprising instructions to:

- provide a storage with a knowledge base, being a SNOMED knowledge base, in a web ontology format;
- receive procedural data representing clinical procedures for evaluation of a patient's health state;
- map the received procedural data in a set of SNOMED expressions;
- convert the SNOMED expressions into statements in the web ontology format;
- upgrade the SNOMED knowledge base with the received procedural data by adding the statements in the SNOMED knowledge base for providing a processable

file with an upgraded version of the SNOMED knowledge base, wherein the knowledge base comprises a graph-based decision tree; and

generate a virtual representation of the decision tree processing.

18. The non-transitory computer readable storage medium according to claim **17**, wherein the instructions comprise use of the upgraded SNOMED knowledge base for classification of a patient's health state by loading the provided processable file in an ontology reader, wherein the ontology reader applies a classification algorithm.

19. The non-transitory computer readable storage medium according to claim **17**, wherein the instructions comprise application of a specific patient instance to the processable file for inference of the patient's health state by an ontology reader for applying a reasoning algorithm.

20. The non-transitory computer readable storage medium according to claim **17**, wherein the processable file provides result data even when a parameterization is incomplete and/or parameters to be processed are inconsistent.

* * * * *