General Artificial Intelligence Method and Computer System

Technical Field

The present invention relates to a method of artificial general intelligence, and to the computer executable steps, computer readable code, one or more computer readable mediums, and one or more computers adapted for such artificial general intelligence. The primary application of the invention is to the field of semantic networks, as may be used as part of an artificial general intelligence or robotics system. However, it will be appreciated that the invention is not limited to this particular field of use.

Background

Associated technologies required for artificial general intelligence and semantic networks construction have already entered the commercial market, such as the relational dependency parsers RelEx, the Stanford Parser, and MiniPar - some of which have a primary focus on semantics extraction. Competing implementations of artificial general intelligence have already been produced, such as IBM Watson (otherwise known as Deep QA), which allows real-time Question-Answer feedback based upon a trained "context" of textual information. For example such systems enable a section of text (which may be generated based upon audio/visual input) to be stored in a database in such a way that detailed queries regarding one or more sections of that information base may be performed. The text often used for this purpose is open license encyclopaedic text (for example Wikipedia).

A number of prior art patents describe artificial general intelligence methods based upon semantic networks construction.

This artificial general intelligence method primarily comprises the breaking down or reduction of language (or its pre-extracted inter-word dependency relations) into a simplistic yet comprehensive set of semantic relationship types, and may be represented as a semantic network. The advantage of seeking a minimal (yet comprehensive) set of rules in semantic relationship construction is that it enables a broader range of queries to be defined in order to obtain (query) a given relationship by eliminating all dependency upon the syntactic structure of the original text (while minimizing “information degradation” - the loss of essential information - in this case otherwise extractable semantic relationships).

The present invention aims to provide a method which is “data rich” and extremely efficient, and therefore fast in execution.

Ideally the general intelligence method should operate independent of the language, the information context (or world), and the way in which the information context has been defined (syntactic structure). It should generally be implemented by generating semantic relationships based upon an information "context" of text (with the subsequent formation of a semantic network), allowing then the querying of this information "context" regarding one or more particular semantic relationships between the concepts (or entities) contained therein.

This artificial general intelligence method may be interpreted as an implementation of Aristotle's Categories, in that it appears to conform somewhat to the divisions presented by Aristotle, as shown in the table below. It is also evident that the semantic network layout described here in distinctly relies on the concepts; 'property' and 'definition' - or "having" and "being", relationships which happen to be used in the construction of the two primary auxiliary verbs ("having" and "being") in English and other (Indo-)European languages.

Aristotelian Category	**Relevant Entity**	**Example**
Primary Substance	Concept Entity without instantiation(s) and which does not define any concept entities (i.e. a named individual)	Socrates (the person)
Secondary Substance	a) Concept Entity or b) Property Entity which is an instantiation of a Concept Entity	a) persons b) a person
Quantity	Quantity/Measure (Property) Entity	five meters tall
Quality	Quality (Property) Entity	happy
Relation	Condition Entity	longer than
Place	Condition Entity	**in** the city
Time	(Time) Condition Entity	yesterday / before the bells ring
Position	Action Entity {Alternate interpretation: Condition Entity}	sitting {Alternate interpretation: **near** the house}
State	(incoming) Action Entity {Alternate interpretation: Property Entity}	armed {Alternate interpretation: having arms}
Action	Action Entity	rowed
Affection	(incoming) Action Entity	to be lanced

It is to be understood that, if any prior art information is referred to herein, such reference does not constitute an admission that the information forms part of the common general knowledge in the art, in Australia or any other country.

Summary of Invention

According to one aspect of the present invention, there is provided a method for artificial intelligence, to be performed on or with the aid of one or more computers, comprising the parsing of textual data, and the extraction of semantic relationships, containing one or more entities.

Preferably, the textual data is extracted from an information context.

Preferably, the information context comprises audio data, obtained through speech recognition.

Preferably, the information context comprises visual data, obtained through object recognition.

Preferably, the information context comprises data obtained through a database or network.

Preferably, the information context comprises computer generated data.

Preferably, the set of semantic relationships are represented by a semantic network.

Preferably, the semantic network comprises concept nodes, property nodes, action nodes, and condition nodes.

Preferably, the semantic network comprises concept connections, property connections, action connections, and condition connections.

Preferably, the concept connections represent definitions, the property connections represent properties, the action connections represent actions, and the condition connections represent conditions.

Preferably, the semantic relationships comprise concepts, properties, actions, and conditions.

Preferably, an entity may have a property.

Preferably, a property is a quality

Preferably, a quality is generated based upon a attributive adjectival modifier.

Preferably, a quality is generated based upon an adverbial modifier.

Preferably, a quality is generated based upon a predicative adjectival modifier.

Preferably, a quality is generated based upon an 'is' statement.

Preferably, a property is a subset or part of the entity.

Preferably, a property is generated based upon a noun.

Preferably, a property is generated based upon a 'has' statement.

Preferably, a property is generated based upon a prenominal modifier of a noun

Preferably, a property is a possession

Preferably, a possession is generated based upon a possessive or genitive modifier of a noun.

Preferably, a property is an entity.

Preferably, a property is a property entity.

Preferably, a property is an instance of a concept.

Preferably, a new property is not assigned for an entity if the property has already been declared in the textual context (e.g. in the given paragraph).

Preferably, an entity may have a definition.

Preferably, a definition is a concept.

Preferably, a definition is an entity or concept entity.

Preferably, each concept entity is a unique concept, either having or not having a representation (or instantiation) in the context.

Preferably, the context is a real or virtual world.

Preferably, the concept entities are stored in an indexed data structure thereby providing fast look up of known semantic relationships based upon concepts (or words).

Preferably, a concept entity may define another concept entity.

Preferably, a concept is generated based upon a noun.

Preferably, a concept is generated based upon an 'is' statement.

Preferably, a concept is generated based upon an appositive of a noun (appositional modifier)

Preferably, an entity may have a condition.

Preferably, a condition is generated based upon a preposition.

Preferably, a condition is generated based upon a relation.

Preferably, a relation is a qualitative relationship between two entities, such as a length comparison.

Preferably, a condition is an entity.

Preferably, a condition has a condition type.

Preferably, a condition type is an entity.

Preferably, a condition type is a condition entity.

Preferably, a condition type is defined by an entity.

Preferably, an condition type is an instance of a concept.

Preferably, a condition is a relative or absolute time or date

Preferably, an entity may have an action.

Preferably, an action is generated based upon a verb.

Preferably, an action is applied to an entity.

Preferably, an action has an action type.

Preferably, an action type is an entity.

Preferably, an action type is an action entity.

Preferably, an action type is defined by an entity.

Preferably, an action type is an instance of a concept.

Preferably, an entity may have a quantity.

Preferably, an entity may have a measure.

Preferably, a measure is a measure of distance, size, time, or repetition.

Preferably, an entity may have a relation.

Preferably, a relation is a qualitative relationship between two entities, such as a length comparison.

Preferably, a query may be performed of the semantic relationships.

According to one aspect of the present invention, there is provided computer executable steps for artificial intelligence, to be performed on or with the aid of one or more computers, comprising the parsing of textual data, and the extraction of semantic relationships, containing one or more entities.

Preferably, the textual data is extracted from an information context.

Preferably, the information context comprises audio data, obtained through speech recognition.

Preferably, the information context comprises visual data, obtained through object recognition.

Preferably, the information context comprises data obtained through a database or network.

Preferably, the information context comprises computer generated data.

Preferably, the set of semantic relationships are represented by a semantic network.

Preferably, the semantic network comprises concept nodes, property nodes, action nodes, and condition nodes.

Preferably, the semantic network comprises concept connections, property connections, action connections, and condition connections.

Preferably, the concept connections represent definitions, the property connections represent properties, the action connections represent actions, and the condition connections represent conditions.

Preferably, the semantic relationships comprise concepts, properties, actions, and conditions.

Preferably, an entity may have a property.

Preferably, a property is a quality

Preferably, a quality is generated based upon a attributive adjectival modifier.

Preferably, a quality is generated based upon an adverbial modifier.

Preferably, a quality is generated based upon a predicative adjectival modifier.

Preferably, a quality is generated based upon an 'is' statement.

Preferably, a property is a subset or part of the entity.

Preferably, a property is generated based upon a noun.

Preferably, a property is generated based upon a 'has' statement.

Preferably, a property is generated based upon a prenominal modifier of a noun

Preferably, a property is a possession

Preferably, a possession is generated based upon a possessive or genitive modifier of a noun.

Preferably, a property is an entity.

Preferably, a property is a property entity.

Preferably, a property is an instance of a concept.

Preferably, a new property is not assigned for an entity if the property has already been declared in the textual context (e.g. in the given paragraph).

Preferably, an entity may have a definition.

Preferably, a definition is a concept.

Preferably, a definition is an entity or concept entity.

Preferably, each concept entity is a unique concept, either having or not having a representation (or instantiation) in the context.

Preferably, the context is a real or virtual world.

Preferably, the concept entities are stored in an indexed data structure thereby providing fast look up of known semantic relationships based upon concepts (or words).

Preferably, a concept entity may define another concept entity.

Preferably, a concept is generated based upon a noun.

Preferably, a concept is generated based upon an 'is' statement.

Preferably, a concept is generated based upon an appositive of a noun (appositional modifier)

Preferably, an entity may have a condition.

Preferably, a condition is generated based upon a preposition.

Preferably, a condition is generated based upon a relation.

Preferably, a relation is a qualitative relationship between two entities, such as a length comparison.

Preferably, a condition is an entity.

Preferably, a condition has a condition type.

Preferably, a condition type is an entity.

Preferably, a condition type is a condition entity.

Preferably, a condition type is defined by an entity.

Preferably, an condition type is an instance of a concept.

Preferably, a condition is a relative or absolute time or date

Preferably, an entity may have an action.

Preferably, an action is generated based upon a verb.

Preferably, an action is applied to an entity.

Preferably, an action has an action type.

Preferably, an action type is an entity.

Preferably, an action type is an action entity.

Preferably, an action type is defined by an entity.

Preferably, an action type is an instance of a concept.

Preferably, an entity may have a quantity.

Preferably, an entity may have a measure.

Preferably, a measure is a measure of distance, size, time, or repetition.

Preferably, an entity may have a relation.

Preferably, a relation is a qualitative relationship between two entities.

Preferably, a query may be performed of the semantic relationships.

Brief Description of Drawings

Fig. 1 shows the generation of an information "context"; text being extracted based upon an audio-visual sequence and network database, from which its language is parsed, dependency relations are defined, semantic relationships are extracted, a semantic network is constructed, then queried,

Fig. 2 shows a basic flow diagram of the semantic relationship generation method according to the present invention,

Fig. 3 shows a detailed flow diagram of the semantic relationship generation method according to the present invention,

Fig. 4 shows the creation of a concept entity ("bike") and property entity node ("wheel"), along with a property definition (concept) entity node ("wheel") based upon a given input text ("Bikes have wheels"). Fig. 4 also shows the creation of concept entity nodes based upon a given input text ("Bikes are machines"),

Fig. 5 shows the creation of property entity nodes ("bike", "wheel", "green"), one of which is a quality ("green"), and their corresponding property definition (concept) entities nodes based upon a given input text ("The bike's wheels are green", or alternatively, "The bike has wheels. The bike's wheels are green"),

Fig. 6 shows the creation of an action node ("ride"), an action definition entity node, a subject concept entity node ("Tom"), and an object property entity node ("bike"), based upon a given input text ("Tom rides the bike quickly"),

Fig. 7 shows the creation of a condition node ("to") between an action node ("go") and a property node ("park") based upon a given input text ("He goes to the park"),

Fig. 8 shows a semantic network with reference linking (pronoun replacement) between entities ("he" with "Tom") within a given textual context (e.g. paragraph) based upon the combination of the input text from Fig. 4, Fig. 5, Fig. 6 and Fig. 7 ("Tom rides the bike quickly. He goes to the park. The bike's wheels are green. bikes are machines"),

Fig. 9 shows the formation of various semantic relationships based upon syntactic prototypes,

Fig. 10 shows a semantic network, with entity references stored in an array (map) of ascending alphabetical order - providing fast look up of concepts, based upon the semantic network shown in Fig. 8,

Fig. 11 shows a semantic network containing one or more concept nodes, property nodes, action nodes, and condition nodes (extract from Wikipedia article "Neuroscience of free will")

Fig. 12 shows a semantic network containing a 'how' query and an existing knowledge base upon which the query is applied ("How did the disaster happen?" and "The disaster happened over night." respectively. The answer given is "night", with condition "over". An example of the answer with a generated textual context is "disaster happen over a night"),

Fig. 13 shows a semantic network containing a 'how much' query and an existing knowledge base upon which the query is applied ("How much milk should Tom pour?" and "Tom should pour 4 litres." respectively. The answer given is "litre", with quantity "4". An example of the answer with a generated textual context is "pour 4 litre"),

Fig. 14 shows a semantic network containing a 'what' query and an existing knowledge base upon which the query is applied ("What is the time?" and "The time is 06:45." respectively. The answer given is "06:45". An example of the answer with a generated textual context is "time is 06:45"),

Fig. 15 shows a semantic network containing a 'when' query and an existing knowledge base upon which the query is applied ("When should they leave?" and "They should leave at six." respectively. The answer given is "six", with condition "at". An example of the answer with a generated textual context is "leave at 6 pm")

Fig. 16 shows a semantic network containing a 'where' query and an existing knowledge base upon which the query is applied ("Where is the ball?" and "The ball is in the park." respectively. The answer given is "park", with condition "in". An example of the answer with a generated textual context is "ball in park"),

Fig. 17 shows a semantic network containing a 'which' query and an existing knowledge base upon which the query is applied ("Which house did Jane buy?" and "Jane bought the blue house." respectively. The answer given is "house". An example of the answer with a generated textual context is "buy house", which has a property blue),

Fig. 18 shows a semantic network containing a 'who' query and an existing knowledge base upon which the query is applied ("Who is that?" and "That is Jim." respectively. The answer given is "Jim". An example of the answer with a generated textual context is "that is a Jim"),

Fig. 19 shows a semantic network containing a 'why' query and an existing knowledge base upon which the query is applied ("Why does the star fall?" and "The star falls because it rains." respectively. The answer given is "rain", with condition "because". An example of the answer with a generated textual context is "fall because rain"),

Fig. 20 shows an example of a query semantic network for application to the semantic network shown in Fig. 8, as well as the subsequent trace and identification of the query variable "_$qVar" (The query provided is "Where does Tom go?". The answer given is "park"),

Fig. 21 shows a different implementation of the semantic network construction based upon the semantic network shown in Fig. 8, in which both Condition and Action entity nodes have been removed, and the connections between their "subjects" and "objects" have instead been named accordingly.

Figure Key - Semantic Network

Entity Nodes:

- White = Concept Entity (White)
- Cyan = Property Entity (Cyan)
    - Aqua = Quality (Property) Entity (Aqua)
    - Grey = Measure (Property) Entity (Light Grey)
    - Purple = Quantity (Property) Entity (Purple)
- Green = Action Entity (Green)
- Red = [Time] Condition Entity (Red)
- Yellow = Query/Answer Entity (Yellow)
- Yellow = Query Traced Entity (Dark Grey)
- Bold = Plural / Progressive Time Condition Entity (**Bold**)

Entity Connections:

- Blue = Definition (Blue)
- Magenta = Instance (Magenta)
- Cyan = Property (Cyan)
- Green = Action (Green)
- Red = Condition (Red)

Fig. 9 Key - Syntactic Prototypes

Formal description	Example Sentence	Example Syntactic Dependency Relation (RelEx)
attributive adjectival modifier	The broken pencil fell apart.	_amod(pencil, broken)
adverbial modifier	Tom runs quickly.	_advmod(run, quickly)
appositive of a noun (appositional modifier)	The fish, a carp, swam deeply.	_appo(fish, carp)
day of month / year modifier	The battle happened on March 11th, 1973.	_date_day(March, 11th) / _date_year(March, 1973)
expletive/filler subject	There is a new movie we can see.	_expl(be, there)
indirect object	The officer gave the youth a ride.	_iobj(give, youth)
unit of distance measure	The rabbit is 20 meters away.	_measure_distance(away, meters)
unit of repetition.	Robert ate 4 times a day.	_measure_per(times, day)
unit of size measure	The girl is 6 feet tall.	_measure_size(tall, feet)
unit of time measure.	The undergraduate student is 18 years old.	_measure_time(old, years)
modifier of a noun	Hamish smoked at the toy shop.	_nn(shop, toy)
direct object, also used for passive nominal subject.	Alice wrote the letter.	_obj(wrote, letter)
parataxis	The guy, Akari said, left early in the morning.	_parataxis(leave, say)
subject of preposition / object of preposition	The bridge is next to the city.	_psubj(next_to, garage) /_pobj(next_to, house)
possessive or genitive modifier of a noun (gen)	Ms. Savata's hand slipped.	_poss(hand, Ms._Savata)
predicative adjectival modifier	Kane is late.	_predadj(Kane, late)
numeric modifier	Margaret brought five pounds.	_quantity(pound, five)
quantity modifier	The punter won almost $1000.	_quantity_mod($, almost)
quantity multiplier	Elexis owes seven hundred dollars.	_quantity_mult(hundred, seven)
subject of a verb	She is the one that broke it.	_subj(be, she) _subj(do, one)
implied preposition "that"	Moses knew I was angry.	_that(know, angry)
adjectival complement (acomp)	The pastry tasted awesome.	_to-be(taste, awesome)
clausal complement (ccomp/xcomp)	Jezel likes to draw.	_to-do(like, draw)

Fig. 11 Key - Wikipedia Extract (Neuroscience of free will)

A recent study by Masao Matsuhashi and Mark Hallett claims to have replicated Libet's findings without relying on subjective report or clock memorization on the part of participants. The authors believe that their method can identify the time at which a subject becomes aware of her own movement. Matsuhashi and Hallet argue that this time not only varies, but often occurs after early phases of movement genesis have already begun (as measured by the readiness potential). They conclude that a person's awareness cannot be the cause of movement, and may instead only notice the movement.

Matsuhashi and Hallett's study can be summarized thus. The researchers hypothesized that, if our conscious intention are what causes movement genesis, the start of an action, then naturally, our conscious intention should always occur before any movement has begun. This is because otherwise, if we ever become aware of a movement only after it has already been started, our awareness could not have been the cause of that particular movement. Simply put, conscious intention must precede action if it is its cause.

To test this hypothesis, Matsuhashi and Hallet had volunteers perform brisk finger movements at random intervals, while not counting or planning when to make such movements, but rather immediately making a movement as soon as they thought about it. An externally controlled stop-signal sound was played at pseudo random intervals, and the volunteers had to cancel their intent to move if they heard a signal while being aware of their own immediate intention to move. Whenever there was an action or finger movement, the authors documented and graphed any tones that occurred before that action. The graph of tones before actions therefore only shows tones before the subject is even aware of their movement genesis, or else they would have stopped or vetoed the movement, and after it is too late to veto the action. This second set of graphed tones is of little importance here.

In this work, "movement gensis" is defined as the brain process of making movement, of which physiological observations have been made via electrodes indicating that it may occur before conscious awareness of intent to move.

By looking to see when tones started preventing actions, the researchers supposedly know the length of time that exists between when a subject holds a conscious intention to move and performs the action of movement. This moment of awareness as seen in the graph below is dubbed T, the mean time of conscious intention to move. It can be found by looking at the border between tones and no tones. This enables the researchers to estimate the timing of the conscious intention to move without relying on the subject's knowledge or demanding them to focus on a clock. The last step of the experiment is to compare time T for each subject with their Event-related potential measures, e.g. seen in this page's lead image, which reveal when their finger movement genesis first begins.

The researchers found that the time of the conscious intention to move T normally occurred too late to be the cause of movement genesis. See the example of a subject's graph below on the right. Although it is not shown on the graph, the subject's readiness potentials tells us that his actions start at -2.8 seconds, and yet this is substantially earlier than his conscious intention to move, time "T", -1.8 seconds. Matsuhashi and Hallet concluded that the feeling of the conscious intention to move does not cause movement genesis; both the feeling of intention and the movement itself are the result of unconscious processing.

Best Mode for Carrying Out the Invention

Referring to Fig. 1, an artificial general intelligence method according to a first embodiment of the present invention will be described. Text is extracted from an information "context", based upon a variety of methods, from a variety of sources, including an audio-visual sequence and network database. Computer generation of text from a virtual world is also shown to highlight the flexibility in implementation of this method. The language of the text is then parsed using a language parser, whereby syntactic dependency relations are extracted. Semantic relationships are then generated. A semantic network is constructed, and then queried using additional input (such as from a user).

The language parser (or preprocessor) is not detailed in this invention, and it may involve the combination of a large number of functions. In this embodiment RelEx is used to perform syntactic dependency relation extraction, and depends upon a number of other software packages. The language parser may include sentence detection (e.g. OpenNLP), spell checking (e.g. Link Grammar Parser), morphological analysis (e.g. Link Grammar Parser / WordNet), entity detection (e.g. GATE), link parsing (e.g. CMU Link Grammar Parser), dependency relation extraction (e.g. RelEx), anaphora resolution (e.g. RelEx), and frame processing (e.g. RelEx). The dependency grammar parser is used to extract the base syntactic dependencies (Carnegie Mellon University's Link Grammer Parser, Sleater and Temperley 1993), and morphological analysis is commonly performed using WordNet (Fellbaum 1998). Morphological functions are used to link entities together that vary only based upon their grammatical context (e.g. tense), and Anaphora resolution refers to reference processing, including pronoun replacement. In this embodiment, reference pre-processing performed by the language parser is not (exclusively) relied upon.

Referring to Fig. 2, a basic flow diagram of the semantic relationship generation method is shown according to the present invention. The syntactic dependency relations used for the creation of semantic network nodes and connections thereof are defined for this embodiment of the invention. The figure shows a basic flow diagram of the semantic relationship generation method. These include a broad class of syntactic relations (such as adjectives, objects/subjects, verbs, etc), and a broad class of semantic relations. The semantic relations defined here are expressed using a number of specific Entity node classes, and connections or "Links" between these nodes. The combination of the nodes and their connections may be represented by a semantic network.

The entity nodes described here include Concept entities, Property entities, Actions entities, and Condition entities. These entity classes are provided examples of their likely mapping to language structures, wherein nouns are commonly identified as Concept entities, adjectives are commonly identified as Property entities, verbs are commonly identified as Action entities, and prepositions are commonly identified as Condition entities. The entity node connections include here Concept Definitions (concept-concept), Property Links (entity-property), Action Links (entity-action-entity), and Condition Links (entity-condition-entity), along with Property Definitions (property-concept), Action Definitions (action-concept) and Condition Definitions (condition-concept).

It may be noted that in this embodiment of the invention, both nouns and adjectives are commonly identified as property entities, wherein for example both the possessive modifier and attributive adjectival modifier syntactic dependency relations are used in the definition of Property Links.

Referring to Fig. 3, a detailed flow diagram of the semantic relationship generation method is shown according to the present invention. The execution flow of the semantic relationship generator used in this embodiment of the invention is outlined, based on the previously extracted syntactical dependency relations. Limited textual examples are provided of relevant syntactic structures. This diagram demonstrates execution flow which is repeated for each sentence of input text.

As indicated in Fig. 2, various combinations of semantic relationships are possible. Referring to Fig. 4, the creation of a concept entity 1 ("bike") and property entity node 2 ("wheel"), along with a property definition (concept) entity node 3 ("wheel") based upon a given input text ("Bikes have wheels") are shown. Fig. 4 also shows the creation of concept entity nodes 4 based upon a given input text ("Bikes are machines"). Referring to Fig. 5, the creation of property entity nodes 5 ("bike", "wheel", "green"), one of which is a quality 6 ("green"), and their corresponding property definition (concept) entities 7 nodes based upon a given input text ("The bike's wheels are green", or alternatively, "The bike has wheels. The bike's wheels are green") are shown. Referring to Fig. 6, the creation of an action node 8 ("ride"), an action definition entity node 9, a subject concept entity node 10 ("Tom"), and an object property entity node 11 ("bike"), based upon a given input text ("Tom rides the bike quickly") are shown. Referring to Fig. 7, the creation of a condition node 12 ("to") between an action node 13 ("go") and a property node 14 ("park") based upon a given input text ("He goes to the park") is shown.

Common entities may also be referred to in the text using different words, such a pronouns, in which case reference linking may be required. Referring to Fig. 8, a semantic network with reference linking (pronoun replacement) between entities 15 ("he" with "Tom") within a given textual context (e.g. paragraph) is shown, based upon the combination of the input text from Fig. 4, Fig. 5, Fig. 6 and Fig. 7 ("Tom rides the bike quickly. He goes to the park. The bike's wheels are green. bikes are machines").

A large number of syntactic prototypes are supported by the present embodiment of the invention. Referring to Fig. 9, the formation of various semantic relationships based upon syntactic prototypes are shown (based on the RelEx standard - see section Brief Description of Drawings).

According to the present embodiment of the invention, the semantic network database is optimised, such that concepts may be searched by name (e.g. during the query semantic network matching process, described below). Referring to Fig. 10, a semantic network, with entity references stored in an array 16 (map) of ascending alphabetical order is shown - providing fast look up of concepts, based upon the semantic network shown in Fig. 8 (i.e, the combination of input text from Fig. 4, Fig. 5, Fig. 6 and Fig. 7).

Referring to Fig. 11, a semantic network containing one or more concept nodes, property nodes, action nodes, and condition nodes (extract from Wikipedia article "Neuroscience of free will" - see section Brief Description of Drawings) is shown. Generated semantic networks may be considered to be of arbitrary size and connectivity, for example, a single concept entity may have a large number of instances (eg property entities, action entities, or condition entities), and an entity may have a large number of properties, actions, or conditions.

The semantic relationships and corresponding semantic network generated supports a variety of user queries (or questions). According to this embodiment of the present invention queries are performed by first translating the query text into an independent set of semantic relationships (and correspondingly forming a independent semantic network, or "query semantic network"), and matching the existing semantic network with the query semantic network using an matching algorithm.

Based upon the accuracy of the match(es), a confidence value is assigned, where in this embodiment the confidence value is a function of the number of nodes matched. The generated query semantic network will often (if not always) have a specific query node specified, indicative of the position of the expected answer node in the original (non-query) semantic network. The matching of such a query node during the semantic network matching process can therefore be considered a critical requirement or per-requisite of a successful match by the algorithm. There may however exist cases of querying where a constructed query semantic network has no specific "query node", in which case an exact answer node cannot be found (i.e. a query node designated in the query semantic network directly corresponding to a relevant answer node in the semantic network). In the present embodiment, queries identified by "which" have no query node designated, as shown for example in Fig. 17, where as most other basic queries have specific nodes assigned (specified by the RelEx standardised "_$qVar" 17, as shown for example in Fig. 12).

Referring to Fig. 12, a semantic network containing a 'how' query and an existing knowledge base upon which the query is applied ("How did the disaster happen?" and "The disaster happened over night." respectively. The answer given is "night" 18, with condition "over" 19. An example of the answer with a generated textual context is "disaster happen over a night") is shown.

Referring to Fig. 13, a semantic network containing a 'how much' query and an existing knowledge base upon which the query is applied ("How much milk should Tom pour?" and "Tom should pour 4 litres." respectively. The answer given is "litre" 20, with quantity "4" 21. An example of the answer with a generated textual context is "pour 4 litre") is shown.

Referring to Fig. 14, a semantic network containing a 'what' query and an existing knowledge base upon which the query is applied ("What is the time?" and "The time is 6:45." respectively. The answer given is "6:45" 22. An example of the answer with a generated textual context is "time is 6:45") is shown.

Referring to Fig. 15, a semantic network containing a 'when' query and an existing knowledge base upon which the query is applied ("When should they leave?" and "They should leave at 6pm." respectively. The answer given is "six" 23, with condition "at" 24. An example of the answer with a generated textual context is "leave at 6 pm") is shown.

Referring to Fig. 16, a semantic network containing a 'where' query and an existing knowledge base upon which the query is applied ("Where is the ball?" and "The ball is in the park." respectively. The answer given is "park" 25, with condition "in" 26. An example of the answer with a generated textual context is "ball in park") is shown.

Referring to Fig. 17, a semantic network containing a 'which' query and an existing knowledge base upon which the query is applied ("Which house did Jane buy?" and "Jane bought the blue house." respectively. The answer given is "house" 27. An example of the answer with a generated textual context is "buy house", which has a property blue 28) is shown.

Referring to Fig. 18, a semantic network containing a 'who' query and an existing knowledge base upon which the query is applied ("Who is that?" and "That is Jim." respectively. The answer given is "Jim" 29. An example of the answer with a generated textual context is "that is a Jim") is shown.

Referring to Fig. 19, a semantic network containing a 'why' query and an existing knowledge base upon which the query is applied ("Why does the star fall?" and "The star falls because it rains." respectively. The answer given is "rain" 30, with condition "because" 31. An example of the answer with a generated textual context is "fall because rain") is shown.

in this embodiment of the present invention, the matching algorithm selected is arbitrary, or implementation dependent. One such matching algorithm involves the following: For every concept node in the query semantic network, finding (searching for / performing an index look up of) the corresponding concept node in the original (non-query) semantic network. Then tracing all connections (instances) from the matched concept entities, where matching connection types and entity node names are found. For every match found during the trace, the confidence value for the trace is incremented. The semantic network trace which gives the highest confidence value (perhaps subject also to the successful match of an entity node corresponding to the query variable) is then taken as corresponding to (or containing) the query's most likely answer - identified within a specific context of the original semantic network. Referring to Fig. 20, an example of a query semantic network is shown, for application to the semantic network shown in Fig. 8, as well as the subsequent trace and identification of the query variable "_$qVar" 32 (The query provided is "Where does Tom go?". The answer given is "park". An example of the answer with a generated textual context is "go to park").

In this embodiment of the present invention, the use/presentation of entity nodes in semantic network construction of the same connectivity (semantic relationships) may be considered arbitrary, or implementation dependent. For example, the use/presentation of condition and action nodes in this present embodiment is implementation dependent, where they can instead be represented as direct connections (between the "subject" and "object" of the action/condition) with the connection type being defined by the action/condition entity (i.e. in that the connection type name corresponds to the action/condition entity name). Referring to Fig. 21, a different implementation of the semantic network construction based upon the semantic network shown in Fig. 8 (i.e. the combination of the input text from Fig. 4, Fig. 5, Fig. 6 and Fig. 7.) is shown, in which both condition and action entity nodes have been removed, and the connections between their "subjects" and "objects" have instead been named accordingly 33. Similarly, the use/presentation of concept nodes in this present embodiment is implementation dependent. Referring to Fig. 10, a different implementation of the semantic network construction based upon the text used to derive the semantic network shown in Fig. 8 (i.e. the combination of the input text from Fig. 4, Fig. 5, Fig. 6 and Fig. 7.) is shown, in which concept nodes entity nodes have been isolated and stored in an array 16.

In accordance with a second embodiment of the present invention there are provided computer executable steps, these computer executable steps corresponding substantially to the steps in the method as previously described for generating semantic relationships in reference to the first embodiment. Again, the input of the textual data may be a provided by a conventional speech/OCR recognition system or database/network query system (e.g. internet or encyclopaedic archive extraction). The input textual data could also be computer software generated based upon a virtual world or information context.

In accordance with a third embodiment of the present invention there is provided computer readable code for carrying out the computer executable steps as described in reference to the second embodiment, when run on the one or more computers.

In accordance with a fourth embodiment of the present invention there is provided one or more computer readable mediums comprising the computer readable code as described in reference to the third embodiment.

In accordance with a fifth embodiment of the present invention there is provided one or more computers adapted for generating semantic relationships, comprising:

- respective one or more processors;
- respective computer readable memory operatively connected to the respective one or more processors and storing the computer readable code described in reference to the third embodiment;
- respective one or more textual data input systems for receiving the textual data and/or query; and
- an output system for outputting the query results data.

In accordance with this fifth embodiment, a first computer and associated textual input data system could be used to generate the textual input data in formation of the semantic relationships, and a second computer and another associated information input data system could be used to generate the query from one or more other information contexts also containing the information during the later query stage. One of these computers, or a third computer, could be used to execute the query, and generate an output of query results in the form of a rating (e.g. a numerical probability rating) or a more simple “yes/no” indication (e.g. a graphic, an indicator light, or an acoustic “beep”) corresponding to the result of the query. Of course there are many different possibilities, including the first and second computers, and the associated textual input data systems, being a common system, and that common computer also being used to execute the query and thereby generate an output of the above described query result. In another embodiment a computer network or cloud could be employed to perform the generation of the semantic relationships, and/or the generation of the query, and/or the query execution, where the computer network or cloud could be connected to one or more devices for extracting textual input from one or more devices (including but not limited to networks and databases).

The textual data input system employed according to this fifth embodiment may be a conventional speech recognition system, or alternatively may comprise an OCR (optical character recognition) system, or a network (eg internet) or database (eg encyclopaedic) feed/query system. The input textual data could also be computer software generated based upon one or more virtual (or otherwise) information contexts.

Interpretation

Wireless

The invention may be embodied using devices conforming to other network standards and for other applications, including, for example other WLAN standards and other wireless standards. Applications that can be accommodated include IEEE 802.11 wireless LANs and links, and wireless Ethernet.

In the context of this document, the term “wireless” and its derivatives may be used to describe circuits, devices, systems, methods, techniques, communications channels, etc., that may communicate data through the use of modulated electromagnetic radiation through a non-solid medium. The term does not imply that the associated devices do not contain any wires, although in some embodiments they might not. In the context of this document, the term “wired” and its derivatives may be used to describe circuits, devices, systems, methods, techniques, communications channels, etc., that may communicate data through the use of modulated electromagnetic radiation through a solid medium. The term does not imply that the associated devices are coupled by electrically conductive wires.

Processes:

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”, “analysing”, “generating”, “deriving” 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.

Processor:

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 device” or a “computing machine” or a “computing platform” may include one or more processors.

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. The processing system further may include a memory subsystem including main RAM and/or a static RAM, and/or

ROM.

Computer-Readable Medium:

Furthermore, a computer-readable carrier medium may form, or be included in a computer program product. A computer program product can be stored on a computer usable carrier medium, the computer program product comprising a computer readable program means for causing a processor to perform a method as described herein.

Networked or Multiple Processors:

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 client machine in server-client network environment, or as a peer machine in a peer-to-peer or distributed network environment (e.g. cloud). The one or more processors may form 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.

Note that while some diagram(s) only show(s) 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.

Additional Embodiments:

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 are for execution on one or more processors. 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. The computer-readable carrier medium carries computer readable code including a set of instructions that when executed on one or more processors cause a 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.

Carrier Medium:

The software may further be transmitted or received over a network via a network interface device. While the carrier medium is shown in an example 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.

Implementation:

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.

Means For Carrying out a Method or Function

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 processor device, 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.

Coupled

Similarly, it is to be noticed that the term coupled, when used in the claims, should not be interpreted as being limitative 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.

Embodiments:

Reference throughout this specification to “one embodiment” 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” 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.

Similarly it should be appreciated that in the above description of example embodiments of the invention, various features of the invention are sometimes grouped together in a single embodiment, figure, 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 of Specific Embodiments are hereby expressly incorporated into this Detailed Description of Specific Embodiments, with each claim standing on its own as a separate embodiment of this invention.

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 in the art. For example, in the following claims, any of the claimed embodiments can be used in any combination.

Specific Details

In the description provided herein, numerous specific details are set forth. However, it is understood that embodiments of the invention may be practised 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.

Terminology

In describing the preferred embodiment of the invention illustrated in the drawings, specific terminology will be resorted to for the sake of clarity. However, the invention is not intended to be limited to the specific terms so selected, and it is to be understood that each specific term includes all technical equivalents which operate in a similar manner to accomplish a similar technical purpose. Terms such as "forward", "rearward", "radially", "peripherally", "upwardly", "downwardly", and the like are used as words of convenience to provide reference points and are not to be construed as limiting terms.

Definitions

As used herein, unless otherwise specified the use of definitions must be interpreted in their context. For example, "entity" or "entities" referenced herein are defined with respect to language, not with respect to reality (material/non-material). For example, an action herein is an entity, where as in reality, an action is not necessarily (/strictly) an entity.

Different Instances of Objects

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.

Comprising and Including

In the claims which follow and in the preceding description of the invention, except where the context requires otherwise due to express language or necessary implication, the word “comprise” or variations such as “comprises” or “comprising” are used in an inclusive sense, i.e. to specify the presence of the stated features but not to preclude the presence or addition of further features in various embodiments of the invention.

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”.

Scope of Invention

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 fall 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.

Although the invention has been described with reference to specific examples, it will be appreciated by those skilled in the art that the invention may be embodied in many other forms. For example, the term "vehicle" or "vehicular" or other similar term as used herein is inclusive of motor vehicles in general such as passenger automobiles including sports utility vehicles (SUV), buses, trucks, various commercial vehicles, watercraft including a variety of boats and ships, aircraft, and the like, and includes hybrid vehicles, electric vehicles, plug-in hybrid electric vehicles, hydrogen-powered vehicles and other alternative fuel vehicles (e.g. fuels derived from resources other than petroleum). As referred to herein, a hybrid vehicle is a vehicle that has two or more sources of power, for example both gasoline-powered and electric-powered vehicles.

Industrial Applicability

It is apparent from the above, that the arrangements described are applicable to the robotics and manufacturing industries.

Claims

The claims defining the invention are as follows:

1. A method for artificial intelligence, to be performed on or with the aid of one or more computers, comprising the parsing of textual data, and the extraction of semantic relationships, containing one or more entities.

2. The method as claimed in Claim 1 wherein the textual data is extracted from an information context.

3. The method as claimed in Claim 2 wherein the information context comprises audio data, obtained through speech recognition.

4. The method as claimed in Claim 2 wherein the information context comprises visual data, obtained through object recognition.

5. The method as claimed in Claim 2 wherein the information context comprises data obtained through a database or network.

6. The method as claimed in Claim 2 wherein the information context comprises computer generated data.

7. The method as claimed in Claim 1 wherein the set of semantic relationships are represented by a semantic network.

8. The method as claimed in Claim 7 wherein the semantic network comprises concept nodes, property nodes, action nodes, and condition nodes.

9. The method as claimed in Claim 7 wherein the semantic network comprises concept connections, property connections, action connections, and condition connections.

10. The method as claimed in Claim 9 wherein the concept connections represent definitions, the property connections represent properties, the action connections represent actions, and the condition connections represent conditions.

11. The method as claimed in Claim 1 wherein the semantic relationships comprise concepts, properties, actions, and conditions.

12. The method as claimed in Claim 1 wherein an entity may have a property.

13. The method as claimed in Claim 12 wherein a property is a quality

14. The method as claimed in Claim 13 wherein a quality is generated based upon a attributive adjectival modifier.

15. The method as claimed in Claim 13 wherein a quality is generated based upon an adverbial modifier.

16. The method as claimed in Claim 13 wherein a quality is generated based upon a

predicative adjectival modifier.

17. The method as claimed in Claim 13 wherein a quality is generated based upon an 'is' statement.

18. The method as claimed in Claim 12 wherein a property is a subset or part of the entity.

19. The method as claimed in Claim 18 wherein a property is generated based upon a noun.

20. The method as claimed in Claim 18 wherein a property is generated based upon a 'has' statement.

21. The method as claimed in Claim 18 wherein a property is generated based upon a prenominal modifier of a noun

22. The method as claimed in Claim 18 wherein a property is a possession

23. The method as claimed in Claim 22 wherein a possession is generated based upon a possessive or genitive modifier of a noun.

24. The method as claimed in Claim 12 wherein a property is an entity.

25. The method as claimed in Claim 24 wherein a property is a property entity.

26. The method as claimed in Claim 24 wherein a property is an instance of a concept.

27. The method as claimed in Claim 12 wherein a new property is not assigned for an entity if the property has already been declared in the textual context (e.g. in the given paragraph).

28. The method as claimed in Claim 1 wherein an entity may have a definition.

29. The method as claimed in Claim 28 wherein a definition is a concept.

30. The method as claimed in Claim 29 wherein a definition is an entity or concept entity.

31. The method as claimed in Claim 30 wherein each concept entity is a unique concept, either having or not having a representation (or instantiation) in the context.

32. The method as claimed in Claim 31 wherein the context is a real or virtual world.

33. The method as claimed in Claim 30 wherein the concept entities are stored in an indexed data structure thereby providing fast look up of known semantic relationships based upon concepts (or words).

34. The method as claimed in Claim 30 wherein a concept entity may define another concept entity.

35. The method as claimed in Claim 29 wherein a concept is generated based upon a noun.

36. The method as claimed in Claim 29 wherein a concept is generated based upon an 'is' statement.

37. The method as claimed in Claim 29 wherein a concept is generated based upon an

appositive of a noun (appositional modifier)

38. The method as claimed in Claim 1 wherein an entity may have a condition.

39. The method as claimed in Claim 38 wherein a condition is generated based upon a preposition.

40. The method as claimed in Claim 38 wherein a condition is generated based upon a relation.

41. The method as claimed in Claim 40 wherein a relation is a qualitative relationship between two entities, such as a length comparison.

42. The method as claimed in Claim 38 wherein a condition is an entity.

43. The method as claimed in Claim 38 wherein a condition has a condition type.

44. The method as claimed in Claim 43 wherein a condition type is an entity.

45. The method as claimed in Claim 44 wherein a condition type is a condition entity.

46. The method as claimed in Claim 43 wherein a condition type is defined by an entity.

47. The method as claimed in Claim 46 wherein an condition type is an instance of a concept.

48. The method as claimed in Claim 38 wherein a condition is a relative or absolute time or date

49. The method as claimed in Claim 1 wherein an entity may have an action.

50. The method as claimed in Claim 49 wherein an action is generated based upon a verb.

51. The method as claimed in Claim 49 wherein an action is applied to an entity.

52. The method as claimed in Claim 49 wherein an action has an action type.

53. The method as claimed in Claim 52 wherein an action type is an entity.

54. The method as claimed in Claim 53 wherein an action type is an action entity.

55. The method as claimed in Claim 52 wherein an action type is defined by an entity.

56. The method as claimed in Claim 55 wherein an action type is an instance of a concept.

57. The method as claimed in Claim 1 wherein an entity may have a quantity.

58. The method as claimed in Claim 1 wherein an entity may have a measure.

59. The method as claimed in Claim 58 wherein a measure is a measure of distance, size, time, or repetition.

60. The method as claimed in Claim 1 wherein an entity may have a relation.

61. The method as claimed in Claim 60 wherein a relation is a qualitative relationship between two entities, such as a length comparison.

62. The method as claimed in Claim 1 wherein a query may be performed of the semantic relationships.

63. Computer executable steps for artificial intelligence, to be performed on or with the aid of one or more computers, comprising the parsing of textual data, and the extraction of semantic relationships, containing one or more entities.

64. The method as claimed in Claim 63 wherein the textual data is extracted from an information context.

65. The method as claimed in Claim 64 wherein the information context comprises audio data, obtained through speech recognition.

66. The method as claimed in Claim 64 wherein the information context comprises visual data, obtained through object recognition.

67. The method as claimed in Claim 64 wherein the information context comprises data obtained through a database or network.

68. The method as claimed in Claim 64 wherein the information context comprises computer generated data.

69. The method as claimed in Claim 63 wherein the set of semantic relationships are represented by a semantic network.

70. The method as claimed in Claim 69 wherein the semantic network comprises concept nodes, property nodes, action nodes, and condition nodes.

71. The method as claimed in Claim 69 wherein the semantic network comprises concept connections, property connections, action connections, and condition connections.

72. The method as claimed in Claim 71 wherein the concept connections represent definitions, the property connections represent properties, the action connections represent actions, and the condition connections represent conditions.

73. The method as claimed in Claim 63 wherein the semantic relationships comprise concepts, properties, actions, and conditions.

74. The method as claimed in Claim 63 wherein an entity may have a property.

75. The method as claimed in Claim 74 wherein a property is a quality

76. The method as claimed in Claim 75 wherein a quality is generated based upon a attributive adjectival modifier.

77. The method as claimed in Claim 75 wherein a quality is generated based upon an adverbial modifier.

78. The method as claimed in Claim 75 wherein a quality is generated based upon a predicative adjectival modifier.

79. The method as claimed in Claim 75 wherein a quality is generated based upon an 'is' statement.

80. The method as claimed in Claim 74 wherein a property is a subset or part of the entity.

81. The method as claimed in Claim 80 wherein a property is generated based upon a noun.

82. The method as claimed in Claim 80 wherein a property is generated based upon a 'has' statement.

83. The method as claimed in Claim 80 wherein a property is generated based upon a prenominal modifier of a noun

84. The method as claimed in Claim 80 wherein a property is a possession

85. The method as claimed in Claim 84 wherein a possession is generated based upon a possessive or genitive modifier of a noun.

86. The method as claimed in Claim 74 wherein a property is an entity.

87. The method as claimed in Claim 86 wherein a property is a property entity.

88. The method as claimed in Claim 86 wherein a property is an instance of a concept.

89. The method as claimed in Claim 74 wherein a new property is not assigned for an entity if the property has already been declared in the textual context (e.g. in the given paragraph).

90. The method as claimed in Claim 63 wherein an entity may have a definition.

91. The method as claimed in Claim 90 wherein a definition is a concept.

92. The method as claimed in Claim 91 wherein a definition is an entity or concept entity.

93. The method as claimed in Claim 92 wherein each concept entity is a unique concept, either having or not having a representation (or instantiation) in the context.

94. The method as claimed in Claim 93 wherein the context is a real or virtual world.

95. The method as claimed in Claim 92 wherein the concept entities are stored in an indexed data structure thereby providing fast look up of known semantic relationships based upon concepts (or words).

96. The method as claimed in Claim 92 wherein a concept entity may define another concept entity.

97. The method as claimed in Claim 91 wherein a concept is generated based upon a noun.

98. The method as claimed in Claim 91 wherein a concept is generated based upon an 'is' statement.

99. The method as claimed in Claim 91 wherein a concept is generated based upon an appositive of a noun (appositional modifier)

100. The method as claimed in Claim 63 wherein an entity may have a condition.

101. The method as claimed in Claim 100 wherein a condition is generated based upon a preposition.

102. The method as claimed in Claim 100 wherein a condition is generated based upon a relation.

103. The method as claimed in Claim 102 wherein a relation is a qualitative relationship between two entities, such as a length comparison.

104. The method as claimed in Claim 100 wherein a condition is an entity.

105. The method as claimed in Claim 100 wherein a condition has a condition type.

106. The method as claimed in Claim 105 wherein a condition type is an entity.

107. The method as claimed in Claim 106 wherein a condition type is a condition entity.

108. The method as claimed in Claim 105 wherein a condition type is defined by an entity.

109. The method as claimed in Claim 108 wherein an condition type is an instance of a concept.

110. The method as claimed in Claim 100 wherein a condition is a relative or absolute time or date

111. The method as claimed in Claim 63 wherein an entity may have an action.

112. The method as claimed in Claim 111 wherein an action is generated based upon a verb.

113. The method as claimed in Claim 111 wherein an action is applied to an entity.

114. The method as claimed in Claim 111 wherein an action has an action type.

115. The method as claimed in Claim 114 wherein an action type is an entity.

116. The method as claimed in Claim 115 wherein an action type is an action entity.

117. The method as claimed in Claim 114 wherein an action type is defined by an entity.

118. The method as claimed in Claim 117 wherein an action type is an instance of a concept.

119. The method as claimed in Claim 63 wherein an entity may have a quantity.

120. The method as claimed in Claim 63 wherein an entity may have a measure.

121. The method as claimed in Claim 120 wherein a measure is a measure of distance, size, time, or repetition.

122. The method as claimed in Claim 63 wherein an entity may have a relation.

123. The method as claimed in Claim 122 wherein a relation is a qualitative relationship between two entities.

124. The method as claimed in Claim 63 wherein a query may be performed of the semantic relationships.

Abstract

A method of generating semantic relationships is described for the purposes of artificial general intelligence, which is to be performed on or with the aid of one or more computers. The method comprises the parsing of textual data, and the extraction of semantic relationships, containing one or more entities, wherein the semantic relationships comprise concepts, properties, actions, and conditions.