

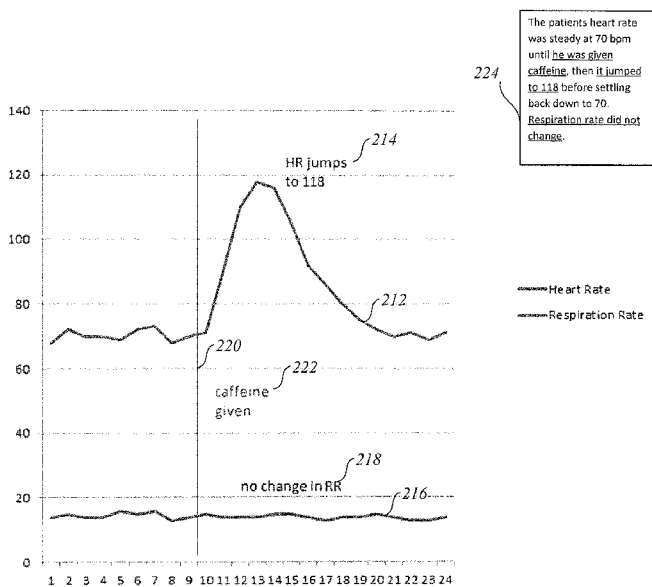


- (51) International Patent Classification:
G06F 17/24 (2006.01)
- (21) International Application Number:
PCT/US2012/053128
- (22) International Filing Date:
30 August 2012 (30.08.2012)
- (25) Filing Language: English
- (26) Publication Language: English
- (71) Applicant (for all designated States except US): DATA2-TEXT LIMITED [GB/GB]; Meston Building, University of Aberdeen, Aberdeen AB24 3UE (GB).
- (72) Inventor; and
- (75) Inventor/Applicant (for US only): REITER, Ehud B. [US/GB]; 43 Don St., Old Aberdeen, Aberdeen AB24 1UH (GB).
- (74) Agents: DOUGLAS, Christopher TL et al.; Alston & Bird LLP, Bank of America Plaza, 101 South Tryon Street, Suite 4000, Charlotte, NC 28280-4000 (US).

- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: METHOD AND APPARATUS FOR ANNOTATING A GRAPHICAL OUTPUT



(57) Abstract: Various methods are provided for generating and annotating a graph. One example method may include determining one or more key patterns in a primary data channel, wherein the primary data channel is derived from raw input data in response to a constraint being satisfied. A method may further include determining one or more significant patterns in one or more related data channels. A method may further include generating a natural language annotation for at least one of the one or more key patterns or the one or more significant patterns. A method may further include generating a graph that is configured to be displayed in a user interface, the graph having at least a portion of the one or more key patterns, the one or more significant patterns and the natural language annotation.

FIG. 2

WO 2014/035403 A1

Declarations under Rule 4.17:

— *of inventorship (Rule 4.17(iv))*

Published:

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

METHOD AND APPARATUS FOR ANNOTATING A GRAPHICAL OUTPUT

TECHNOLOGICAL FIELD

Embodiments of the present invention relate generally to natural language generation technologies and, more particularly, relate to a method, apparatus, and computer program product for textually annotating a graphical output.

5

BACKGROUND

In some examples, a natural language generation (NLG) system is configured to transform raw input data that is expressed in a non-linguistic format into a format that can be expressed linguistically, such as through the use of natural language. For example, raw input data may take the form of a value of a stock market index over time and, as such, the raw input data may include data that is suggestive of a time, a duration, a value and/or the like. Therefore, an NLG system may be configured to input the raw input data and output text that linguistically describes the value of the stock market index. For example, “securities markets rose steadily through most of the morning, before sliding downhill late in the day.”

Data that is input into a NLG system may be provided in, for example, a recurrent formal structure. The recurrent formal structure may comprise a plurality of individual fields and defined relationships between the plurality of individual fields. For example, the input data may be contained in a spreadsheet or database, presented in a tabulated log message or other defined structure, encoded in a ‘knowledge representation’ such as the resource description framework (RDF) triples that make up the Semantic Web and/or the like. In some examples, the data may include numerical content, symbolic content or the like. Symbolic content may include, but is not limited to, alphanumeric and other non-numeric character sequences in any character encoding, used to represent arbitrary elements of information. In some examples, the output of the NLG system is text in a natural language (e.g. English, Japanese or Swahili), but may also be in the form of synthesized speech.

BRIEF SUMMARY

Methods, apparatuses, and computer program products are described herein that are configured to generate a graph that is configured to display one or more key patterns that are detected in a data channel. In some example embodiments, the graph may also include one or more significant patterns in one or more related channels and/or events. In further examples, a time period or duration of the data shown in the graph may be selected such that the displayed graph illustrates the portion of the data channel that contains the one or more key patterns. The output graph is further configured to include textual annotations that provide a textual comment, phrase or otherwise is configured to explain, using text, the one or more key patterns, the one or more significant patterns and/or the events in a contextual channel in natural language. In further examples, the textual annotations are generated from the raw input data and further are designed, in some examples, to textually describe identified patterns, anomalies and/or the context of the graph. In some examples, a narrative may be included with the graph that provides situational awareness or an overview of the data/patterns displayed on and/or off of the graph. Advantageously, for example, the graph is configured to visually provide situational awareness to a viewer.

BRIEF DESCRIPTION OF THE DRAWINGS

Having thus described embodiments of the invention in general terms, reference will now be made to the accompanying drawings, which are not necessarily drawn to scale, and wherein:

Figure 1 is a schematic representation of a graphical annotation system that may benefit from some example embodiments of the present invention;

Figure 2 illustrates an example graphical output in accordance with some example embodiments of the present invention;

Figure 3 illustrates a block diagram of an apparatus that embodies a graphical annotation system in accordance with some example embodiments of the present invention; and

Figures 4-6 illustrate flowcharts that may be performed by a graphical annotation system in accordance with some example embodiments of the present invention.

DETAILED DESCRIPTION

Example embodiments will now be described more fully hereinafter with reference to the accompanying drawings, in which some, but not all embodiments are shown. Indeed, the embodiments may take many different forms and should not be construed as limited to the embodiments set forth herein; rather, these embodiments are provided so

that this disclosure will satisfy applicable legal requirements. Like reference numerals refer to like elements throughout. The terms “data,” “content,” “information,” and similar terms may be used interchangeably, according to some example embodiments, to refer to data capable of being transmitted, received, operated on, and/or stored. Moreover, the
5 term “exemplary”, as may be used herein, is not provided to convey any qualitative assessment, but instead merely to convey an illustration of an example. Thus, use of any such terms should not be taken to limit the spirit and scope of embodiments of the present invention. The terms graph or graphical output may be construed to comprise a graph that is configured to be displayable in a user interface, but may also describe an
10 input into a graphing system such that a graph may be created for display in the user interface. As such, the terms graph or graphical output may be used interchangeably herein.

In some example embodiments described herein, the apparatus, method and computer program product is configured to generate a graph having a scale (e.g.
15 amplitude (y-axis) and/or time scale (x-axis)) that advantageously displays one or more data channels (e.g. a first or primary data channel, a secondary or related data channel and/or the like) that are derived from raw input data, one or more natural language text annotations and/or a narrative describing raw input data. As such, advantageously, a user viewing the graph, in a user interface or using other viewing means, may be provided with
20 situational awareness with regard to the patterns shown on the graph as well as the events and or patterns that may have influenced the patterns shown on the graph.

Situational awareness may be defined as the perception of environmental elements with respect to time and/or space, the comprehension of their meaning, and the projection of their status after some variable has changed, such as time, or based on the
25 happening of an event such as an alarm or alert. In other words, situational awareness is a state achieved when information that is qualitatively and quantitatively determined as suitable for particular purpose is made available to a user by engaging them in in an appropriate information exchange pattern or mental model. Situational awareness involves being aware of what is happening in the vicinity of a person or event to
30 understand how information, events, and/or one's own actions may impact goals and objectives, both immediately and in the near future. Situational awareness may also be related to the perception of the environment critical to decision-makers in complex, dynamic areas from aviation, air traffic control, power plant operations, military command and control, engineering, machine monitoring, oil and gas, power plant monitoring,
35 nuclear energy and emergency services such as firefighting and policing. Lacking or inadequate situational awareness has been identified as one of the primary factors in accidents attributed to human error. Accurate mental models are one of the prerequisites

for achieving situational awareness. A mental model can be described as a set of well-defined, highly-organized yet dynamic knowledge structures developed over time from experience.

5 In some examples, a first or primary data channel may be selected for inclusion in a graph based on a selection by a user, via a user interface, may be selected based on the happening of a condition such as, but not limited to, an alert, an alarm, an anomaly, a violation of a constraint, a warning, a predetermined condition and/or the like. Alternatively or additionally the selection of the primary data channel may be determined based on the occurrence and/or detection of a pattern in the primary data channel.

10 In some example embodiments, a secondary or related data channel may also be selected. In some cases, there may be a plurality of secondary or related data channels. The secondary or related data channel may be selected for inclusion in a graph based on the detection of anomalous, unexpected or otherwise flagged behavior in the second or related channel. In some examples, the second or related channel is compared to one or more patterns in the primary data channel over a similar time period. For example, a first data channel may indicate a rise in heart rate, whereas a second data channel may indicate a stable or even a decline in respiration rate. Generally respiration rate rises with heart rate, and, as such, a stable respiration rate is generally unexpected. In some examples, unexpected behavior may lead to a life threatening condition, be indicative of a dangerous condition or the like.

20 Relationships between data channels may be defined as anomalous behavior by a qualitative model such as a domain model. A domain model is a representation of information about the domain. For example a domain model may contain an ontology that specifies the kinds of objects and concepts and the like that may exist in the domain in concrete or abstract form, properties that may be predicated of the objects and concepts and the like, relationships that may hold between the objects concepts and the like, and representations of any specific knowledge that is required to function in the domain. In some example multiple domain models may be provided for a single domain. Example domains may include, but are not limited to, medical, oil and gas, industrial, weather, legal, financial and/or the like. Alternatively or additionally, a plurality of related channels may be included, for example pulse rate, oxygen levels, blood pressure and/or the like.

30 In some examples, patterns (e.g. a trend, spike, step or the like) may be detected or otherwise identified in the primary data channel and/or in the one or more secondary data channels. Once a pattern is detected in the primary data channel and/or the one or more secondary data channels, an importance level or importance is assigned to each of the patterns. In the primary data channel an importance level may be defined based on

thresholds, constraints, predefined conditions or the like. In the secondary data channels an importance level may also be assigned based on thresholds, constraints, predefined conditions or the like, however an importance level may also be assigned based on the relationship between the secondary data channels and the primary data channels and/or the relationships between the patterns detected in the primary data channels and the patterns detected in the secondary data channels. A pattern in the primary channel may be defined as a key pattern in an instance in which the importance level of the pattern exceeds or otherwise satisfies a predefined importance level. Likewise, a significant pattern is a pattern in a secondary data channel that exceeds or otherwise satisfies a predefined importance level. In some examples, a pattern in the one or more secondary channels may also be classified as a significant pattern if it represents an anomaly or otherwise unexpected behavior when compared with the primary data channel.

In some example embodiments, a contextual channel may also be selected. A contextual channel is a data channel that provides a background or circumstance information that may have caused or otherwise influenced the one or more key patterns and/or the one or more significant patterns (e.g. proximate cause). For example, a contextual channel may indicate an event, such as a medical treatment that was applied at the time of or just prior to the rise of the heartbeat and/or the fall or steady state of the respiration rate. Alternatively or additionally, a plurality of data channels may also be selected for inclusion in a graph based on an anomaly or unexpected behavior.

Alternatively or additionally, one or more data channels may be selected for inclusion in a graph even though the one or more data channels are representative of expected behavior. For example, in the medical domain, a medical professional may expect to see both heart rate and respiration rate on a graph even if both are behaving in expected ways, since expected behavior may be indicative of an important result, namely a clean bill of health.

In yet further example embodiments, events may also be generated for display in the graph. An event may be described in a contextual channel, may be entered into an event log that is input with the raw input data or may be inferred. For example, caffeine administration may be entered as an explicit event in a patient record (e.g. in an event log), the caffeine could be detected by a change in one or data channels which record what medication is being administered through an IV line and/or the caffeine administration may be inferred based on a spike in heart rate. In instances in which an event is identified that satisfies an importance threshold, the event may be displayed as a visual annotation. In an example in which a graph is displayed, events may be displayed as a vertical line (see e.g., Fig. 2). Alternatively or additionally events may be generated as a horizontal line with indicators showing the multiple occurrences of an event and/or

the like. In other visualizations, events may be displayed via text, indicator or other visual outputs.

In some example embodiments, a scale may be selected for the graph based on the primary data channel, the secondary data channel or the like. The scale may be
5 determined based on a time period or duration in which a pattern that satisfies an importance threshold is identified, anomalous behavior occurs in a related data channel and/or the like. Alternatively or additionally the time period may be set by a user, may be a time period that is significant or specifically identified on the basis of properties of the domain, or the like.

10 In further example embodiments, textual annotations and/or a narrative may be included with the graph. The textual annotations and/or the narrative may be provided by a natural language generation system that is configured to generate one or more textual annotations in the form of sentences or phrases that describe the patterns in the data channels, expected or unexpected behavior, an event, a contextual channel and/or the
15 like. Additionally, in some examples, the sentences or phrases may take the form of stand-alone text that provides situational awareness and/or situational analysis of the graph. In some examples, situation analysis text may be configured to include pattern descriptions that contribute to narrative coherence, background information or the like. The textual annotations may be located on the graph, such as at the location where the anomalies and/or the patterns are represented in the graph. Alternatively or additionally,
20 the narrative or situational awareness text may be displayed on or near the graph in some examples. Whereas, in other examples, the narrative or situational text may be contained in a separate file or may be generated before/after or otherwise separately from the generation of the graph. Alternatively or additionally, the textual annotations and/or
25 narrative may be provided via speech or other available modalities.

Based on the one or more channels derived from the raw input data, the contextual channel and/or the annotations, the systems and methods described herein are configured to generate a graph for display. The graph is configured to display a time scale that contains those identified sections (e.g. key patterns and/or significant patterns)
30 in the one or more data channels, the textual annotations, additional available visual annotations and/or the like. In some example embodiments, user interaction with the narrative text may result in an annotation on the graphical output to be highlighted. Similarly selection of an annotation may highlight narrative text related to the annotation. Alternatively or additionally, the annotations may include a symbol or other reference
35 numeral that is indicative of or otherwise related to the narrative. For example, the narrative may indicate that a first key pattern is indicated by an arrow, a circle, a box, a reference number or the like in the graph.

Figure 1 is an example block diagram of example components of an example graphical annotation environment 100. In some example embodiments, the graphical annotation environment 100 comprises a data analyzer 102, a data interpreter 104, a graphical annotation engine 106, a natural language generation system 108 and one or more data sources, such as but not limited to, raw input data 110, an event log 112 and a domain model 114. In some example embodiments, historical data may also be accessed and/or otherwise analyzed. The data analyzer 102, a data interpreter 104, graphical annotation engine 106, a natural language generation system 108 make take the form of, for example, a code module, a component, circuitry or the like. The components of the graphical annotation environment 100 are configured to provide various logic (e.g. code, instructions, functions, routines and/or the like) and/or services related to the generation and/or annotation of a graphical output.

In some example embodiments, the data analyzer 102 is configured to input raw data, such as the raw data contained in the raw input data 110. The receipt or input of the raw input data may occur in response to an alarm condition (e.g. an alarm received from a source such as, but not limited to, another system, another monitoring system or the like), a violation of a constraint (e.g. a data value over a threshold, within a threshold for a period of time and/or the like), a user input or the like. Alternatively or additionally the data analyzer 102 may be configured to receive or input raw input data continuously or semi-continuously, such as via a data stream, and determine an importance of the raw input data (e.g., whether the data violates a constraint, satisfies a threshold and/or the like).

Raw input data may include data such as, but not limited to, time series data that captures variations across time (e.g. profits, rainfall amounts, temperature or the like), spatial data that indicates variation across location (e.g. rainfall in different regions), or spatial-temporal data that combines both time series data and spatial data (e.g. rainfall across time in different geographical output areas). The raw input data contained or otherwise made accessible by the raw input data 110 may be provided in the form of numeric values for specific parameters across time and space, but the raw input data may also contain alphanumeric symbols, such as the RDF notation used in the semantic web, or as the content of database fields. The raw input data 110 may be received from a plurality of sources and, as such, data received from each source, sub source or data that is otherwise related may be grouped into or otherwise referred to as a data channel.

The data analyzer 102 is configured to detect patterns and trends in the one or more data channels that are derived from the raw input data to provide a set of abstractions from the raw input data in the data channels. For example, a time-series dataset may contain tens of thousands of individual records describing the temperature at various points on a component piece of machinery over the course of a day with a sample

once every two or three seconds. Trend analysis may then be used to identify that the temperature changes in a characteristic way throughout certain parts of the day. As such, trend analysis is configured to abstract those changes into an abstraction that is representative of the change over time. In some example embodiments, the data
5 analyzer 102 may be configured to fit a piecewise linear model to the data received in the primary data channel, related data channel or the like.

In some example embodiments, the data analyzer 102 is further configured to determine a first or primary data channel. The primary data channel is generally related, for example, to the raw input data and/or the data channel having data values that caused
10 or otherwise related to the alarm condition, a data channel identified by a user action or a data channel that has otherwise been provided to the data analyzer 102. In some example embodiments, the data analyzer 102 may also be configured to identify data channels that are related to the primary data channel. Alternatively or additionally, relations between data channels may be defined by the domain model 114 and input into
15 the data analyzer 102.

The data analyzer 102 may then identify trends, spikes, steps or other patterns in the data channels to generate abstractions that summarize the patterns determined in the primary data channel and/or the other related data channels. Alternatively or additionally, the data analyzer 102 may also be configured to perform pattern detection on the raw
20 input data irrespective of data channels or the receipt of an alarm condition.

A data interpreter, such as data interpreter 104, may then be configured to input the abstractions and determine an importance level and/or relationships between the abstractions identified in the one or more data channels. In order to determine the importance level and relationships, the data interpreter 104 may access the domain
25 model 114 directly or indirectly via the data analyzer 102 or the like. The domain model 114 may contain information related to a particular domain or industry. In some examples, the domain model 114 may provide single data channel limits related to normal behaviors in a domain (e.g. normal ranges), information related to anomalous behaviors and/or the like. In other examples the domain model 114 may describe relationships between
30 various events and/or phenomena in multiple data channels. For example in a weather domain, a domain model may include wind speeds that are related to hurricane type events or temperatures that may cause harm to humans or other animals or may cause damage or interference to shipping. Extreme weather events may be labeled as important, whereas typical temperatures may not be marked as important.

In some example embodiments, the data interpreter 104 may be configured to determine the importance of the one or more detected patterns in the primary data channel, such as by using the domain model 114. The data interpreter 104 may assign an
35

importance level based on the pattern itself (e.g. magnitude, duration, rate of change or the like), defined constraints (e.g. defined thresholds or tolerances), temporal relationships between the pattern in the primary data channel and patterns in other related data channels and/or the like. For example, a heart rate over 170 beats per minute, or 100 mile per hour winds, may be assigned a high level of importance. In some examples, the patterns and/or the constraints may be defined by the domain model 114.

Using the importance level, the data interpreter 104 may assign certain ones of the patterns as key patterns. A key pattern may be selected based on a pre-determined importance level, such as a threshold defined by a user or a constraint defined by the domain model 114. Alternatively or additionally, key patterns may be selected based on those patterns in the primary data channel with the highest level of importance, based on the alarm condition and/or the like. For example any wind readings over 50 miles per hour may be designated as key patterns, whereas in other examples only the highest wind reading over a time period may be a determined to be a key pattern. In other examples, the importance level determination may be performed over a plurality of time scales that may be user defined (e.g., one hour, one day, one week, one month and/or the like).

In some example embodiments, the data interpreter 104 may also be configured to determine the importance of patterns detected in one or more secondary or related data channels. In some examples, the data interpreter 104 may determine one or more patterns in the related data channels that overlap time-wise or occur within the same time period as the patterns in the primary data channel. The data interpreter 104 may then mark the one or more patterns in the related channels as expected, unexpected or as having or not having some other property using the domain model 114. For example, the domain model may suggest that the one or more patterns in the related data channel were expected to rise as they did in the primary channel. By way of example, as winds are rising, a wave height may then be expected to rise. In other cases the behavior of the one or more related channels may be unexpected or may be anomalous when compared to the behavior of the primary data channel.

The data interpreter 104 may be configured to instantiate a plurality of messages based on the raw input data derived from the key events, the significant events, the primary data channel, the one or more related data channels, the historical data, the events (e.g. in the event log 112), the contextual channel and/or the like. In order to determine the one or more messages, the importance level of each of the messages and relationships between the messages, the data interpreter 104 may be configured to access the domain model 104 directly or indirectly via the data analyzer 102 or the like.

In some examples, messages are language independent data structures that correspond to informational elements in a text and/or collect together underlying data in

such a way that the underlying data can be linguistically expressed. In some examples, messages are created based on a requirements analysis as to what is to be communicated for a particular scenario (e.g. for a particular domain). A message typically corresponds to a fact about the underlying data (for example, the existence of some
5 observed event) that could be expressed via a simple sentence (although it may ultimately be realized by some other linguistic means). For example, to linguistically describe wind, a user may want to know a speed, a direction, a time period or the like, but also the user wants to know changes in speed over time, warm or cold fronts, geographic areas and or the like. In some cases, users do not even want to know wind speed, they
10 simply want an indication of a dangerous wind condition. Thus, a message related to wind speed may include fields to be populated by data related to the speed, direction, time period or the like, and may have other fields related to different time points, front information or the like. The mere fact that wind exists may be found in the data, but to linguistically describe "light wind" or "gusts" different data interpretation must be
15 undertaken as is described herein.

The one or more patterns may be marked as significant patterns based on the domain model 114. For example, patterns in the related data channel that have an importance level above a predetermined threshold defined by the domain model 114 may be marked as significant patterns. In some example embodiments, unexpected patterns
20 are also categorized as significant patterns as they are suggestive of a particular condition or fault. Other patterns may be determined to be significant patterns based on one or more constraints on channel value (e.g. expected range of values or the like), data anomalies, patterns marked as neither expected or unexpected that satisfy an importance level, and/or the like.

In further example embodiments, the data interpreter 104 may be configured to determine and/or infer one or more events from the one or more data channels. Events may include specific activities that may influence the one or more key patterns and/or may have caused the one or more significant patterns. In some examples, the one or more events may be inferred based in context with the one or more patterns in the
30 primary data channel. Alternatively or additionally events may be provided as a separate channel, such as a contextual channel, in the raw input data 110 or may be provided directly, such as in an event log 112, to the data interpreter 104.

The one or more key patterns and the one or more significant patterns may be input into the graphical annotation engine 106 and the natural language generation
35 system 108 to enable the generation of a graphical output and/or natural language annotations. In some example embodiments, the graphical annotation engine 106 is configured to generate a graphical output having one or more textual annotations, such

as the graphical output displayed with reference to Figure 2. The graphical output and the one or more textual annotations are configured to be generated by one or more of a scale determination engine 120, an annotation location determiner 122 and a graphical output generator 124.

5 In some example embodiments the scale determination engine 120 is configured to determine a time scale (e.g. x-axis) to be used in the graphical output. The scale determination engine 120 may determine a time period that captures or otherwise includes one or more of the key patterns. In some example embodiments, the time period may be chosen based on the highest number of key patterns, whereas in other
10 embodiments the time scale chosen may include each of the one or more key patterns and/or each of the one or more significant patterns. Alternatively or additionally, the scale determination engine 120 may also determine a scale for the amplitude or y-axis of the graph.

 An annotation location determiner 122 is configured to place one or more
15 annotations, such as textual or visual annotations, on a graphical output produced by the graphical output generator 124. As is described herein, natural language annotations may be generated, such as by the natural language generation system 108, to explain or otherwise describe the one or more key patterns in the primary data channel. In an instance in which the data interpreter 104 determines one or more significant patterns,
20 natural language annotations may also be generated to explain the one or more significant patterns in the related data channels.

 The annotation location determiner 122 is further configured to place an annotation on the graphical output in the proximity of the key pattern or the significant pattern. In other example embodiments, the annotations may otherwise be linked to the
25 graphical output by using reference lines, highlights or other visual indications on or around the graphical output.

 In some example embodiments a natural language generation system, such as natural language generation system 108, is configured to generate phrases, sentences, text or the like which may take the form of natural language annotations. Other linguistic
30 constructs may be generated in some example embodiments. The natural language generation system 108 comprises a document planner 130, a microplanner 132 and/or a realizer 134. Other natural language generation systems may be used in some example embodiments, such as a natural language generation system as described in Building Natural Language Generation Systems by Ehud Reiter and Robert Dale, Cambridge
35 University Press (2000), which is incorporated by reference in its entirety herein.

 The document planner 130 is configured to input the one or more patterns from the data interpreter in the form of messages and determine how to use those messages

to describe the patterns in the one or more data channels derived from the raw input data. The document planner 130 may comprise a content determination process that is configured to select the messages, such as the messages that describe the key patterns and/or the significant patterns, that are to be displayed in the graphical output by the graphical annotation engine 106. In some examples the content determination process may be related to or otherwise limited by the scale determined by the scale determination engine 120. The document planner 130 may also comprise a structuring process that determines the order of messages referring to the key patterns and/or significant patterns to be included in a narrative and/or the natural language annotations.

In some example embodiments, the document planner 130 may access one or more text schemas for the purposes of content determination and document structuring. A text schema is a rule set that defines the order in which a number of messages are to be presented in a document. For example, an event (e.g. medication injection) may be described prior to a key pattern (e.g. rise in heart rate). In other examples, a significant pattern (e.g. falling or steady respiratory rate) may be described after, but in relation to, a description of the key pattern (e.g. rise in heart rate). The output of the document planner 130 may be a tree-structured object or other data structure that is referred to as a document plan. In an instance in which a tree-structured object is chosen for the document plan, the leaf nodes of the tree may contain the messages, and the intermediate nodes of the tree structure object may be configured to indicate how the subordinate nodes are related to each other.

The microplanner 132 is configured to modify the document plan from the document planner 130, such that the document plan may be expressed in natural language. In some example embodiments, the microplanner 132 may perform aggregation, lexicalization and referring expression generation. In some examples, aggregation includes, but is not limited to, determining whether two or more messages can be combined together linguistically to produce a more complex sentence. For example, one or more key patterns may be aggregated so that both of the key patterns can be described by a single sentence. Alternatively or additionally, aggregation may not be performed in some instances so as to enable stand-alone interpretation if a portion of the natural language text is shown as an annotation independently on a graphical output.

In some examples, lexicalization includes, but is not limited to, choosing particular words for the expression of concepts and relations. In some examples, referring expression generation includes, but is not limited to, choosing how to refer to an entity so that it can be unambiguously identified by the reader. The output of the microplanner 132, in some example embodiments, is a tree-structured realization specification whose leaf-

nodes are sentence plans, and whose internal nodes express rhetorical relations between the leaf nodes.

The realizer 134 is configured to traverse the tree-structured realization specification to express the tree-structured realization specification in natural language.

5 The realization process that is applied to each sentence plan makes use of a grammar which specifies the valid syntactic structures in the language and further provides a way of mapping from sentence plans into the corresponding natural language sentences. The output of the process is, in some example embodiments, a well-formed natural language text. In some examples, the natural language text may include embedded mark-up. The
10 output of the realizer 134, in some example embodiments, is the natural language annotations that are configured to be on or in proximity to a graphical output. The realizer may also output situational analysis text or a narrative that is configured to describe or otherwise summarize the one or more key patterns, the one or more significant patterns, the one or more contextual channels, and/or the one or more events to be displayed in
15 the graphical output. Alternatively or additionally, the natural language annotations and/or the narrative may describe data that is not included on the graph to provide additional situational awareness.

The graphical output generator 124 is configured to generate a graphical output within the determined scale. Thus, the graphical output includes the raw input data for the
20 primary data channel and/or any related data channels within the determined scale. The graphical output generator 124 is further configured to display the natural language annotations on the graphical output as well as a narrative that describes the data channels displayed in the graphical output. In some example embodiments, the natural language annotations may be interactively linked to the graphical output. For example,
25 phrases within the narrative may be underlined or otherwise highlighted such that in an instance in which the underlined or otherwise highlighted phrases are selected, a natural language annotation may be shown or otherwise emphasized on the graphical output. Alternatively or additionally, by selecting a natural language annotation on the graphical output, the graphical output generator in conjunction with a user interface may underline
30 or otherwise highlight a corresponding phrase in the narrative. Alternatively or additionally, other visualizations may be provided by the graphical output generator 124 in conjunction with or in replacement of the graph or graphical output, such as, but not limited to, a visual image, a video, a chart and/or the like.

Figure 2 illustrates an example graphical output having multiple data channels in
35 accordance with some example embodiments of the present invention. Figure 2 provides a graphical output that visually represents the behavior of heart rate and respiration rate in response to an application of caffeine over a period of time. The following example

table (e.g. raw input data) illustrates a primary data channel (e.g. heart rate) and a related data channel (e.g. respiration rate):

Time	Heart Rate	Respiration Rate
1	68	14
2	72	15
3	70	14
4	70	14
5	69	16
6	72	15
7	73	16
8	68	13
9	70	14
10	71	15
11	90	14
12	110	14
13	118	14
14	116	15
15	105	15
16	92	14
17	86	13
18	80	14
19	75	14
20	72	15
21	70	14
22	71	13
23	69	13
24	71	14

As is illustrated in the table above, there is a rapid change of heart rate between
5 time point 10 and time point 11 indicating a change in condition. As such, a data analyzer
102, a data interpreter 104 and/or the like may receive an indication of an alarm
condition, may determine that such a spike is representative of an alarm condition, may
receive an indication by the user or the like. In response, the data analyzer 102, the data
10 interpreter 104 or the like may cause the heart rate data channel to be selected as the
primary data channel. In other embodiments, a user, domain knowledge or the like may
indicate that the heart rate channel is selected to be the primary data channel.

In an instance in which a primary data channel is identified, in this case heart rate,
the data analyzer 102, a data interpreter 104 and/or the like may determine whether one
or more key patterns are present in the primary data channel 212. In this example, a key
15 pattern may be the rapid change of heart rate between time point 10 and time point 11,
but it may also include the rise and fall (e.g. spike) of the heart rate between time points

10 and 19, with the high point being at time point 13. In some examples, there may be multiple key patterns in the primary data channel.

The data interpreter 104 may then determine whether there is a secondary or related data channel that contains a significant pattern, such as secondary channel 216 (e.g. respiration rate), that has a significant pattern (e.g. no change when a change is generally expected) in a corresponding time period. As described herein, the corresponding time period may be the same time period or may be a later time period when compared to the time period of the key patterns. Further, the corresponding time period may, in some examples, be defined by a domain model, such as domain model 114.

The data analyzer 102, a data interpreter 104 and/or the like may access data related to the respiration rate of the patient during the same or similar time period. Upon reviewing the corresponding data, the data analyzer 102, the data interpreter 104 and/or the like may determine that the secondary channel 216 was relatively flat and, based on the domain model, such a behavior was unexpected. As described herein, unexpected behavior in a related data channel is a significant pattern. Thus, Figure 2 provides a key pattern, namely the heart rate change between time periods 10 and 19 and a significant pattern represented by the relatively steady respiration rate over the same time period. Figure 2 further comprises an event 220 that was derived from a contextual channel, an event log or the like. The event 220 corresponds to the application of caffeine in this example.

The natural language generation system 108 may input the primary data channel, the secondary data channel, the contextual channel, events, and/or the like. As such, the natural language generation system 108 may be configured to generate a narrative, such as narrative 224, and one or more textual annotations, such as textual annotations 214, 218 and 222. In some examples, the textual annotations are configured to describe the one or more key patterns, the one or more significant patterns, the one or more events, the contextual channel or the like.

The graphical annotation engine 106 may input the key pattern and the significant pattern and may determine a time scale for the graph, such as by the scale determination engine 120. In Figure 2, the scale chosen is configured to highlight the key pattern, the significant pattern and the caffeine event. Based on the scale, the primary and secondary data channels are represented graphically by the graphical output generator 124. The annotation location determiner 122 may locate the key pattern, significant pattern and/or event on the graph and then may assign a location of a textual annotation in a nearby area, such as is shown in Figure 2. One or more text annotations, such as textual

annotations 214, 218 and 222 may then be added to the graph in conjunction with a narrative 224.

Figure 3 is an example block diagram of an example computing device for practicing embodiments of an example graphical annotation environment. In particular, Figure 3 shows a computing system 300 that may be utilized to implement a graphical annotation environment 100 having a data analyzer 102, a data interpreter 104, a graphical annotation engine 106, a natural language generation system 108 and/or a user interface 310. One or more general purpose or special purpose computing systems/devices may be used to implement the data analyzer 102, the data interpreter 104, the graphical annotation engine 106, the natural language generation system 108 and/or the user interface 310. In addition, the computing system 300 may comprise one or more distinct computing systems/devices and may span distributed locations. For example, in some embodiments, the natural language generation system 108 may be accessible remotely via the network 350. In other example embodiments, one or more of the data analyzer 102, the data interpreter 104, the graphical annotation engine 106, the natural language generation system 108 and/or the user interface 310 may be configured to operate remotely. In some example embodiments, a pre-processing module or other module that requires heavy computational load may be configured to perform that computational load and thus may be on a remote device or server. For example, the data analyzer 102 and/or the data interpreter 104 may be accessed remotely. Furthermore, each block shown may represent one or more such blocks as appropriate to a specific example embodiment. In some cases one or more of the blocks may be combined with other blocks. Also, the data analyzer 102, the data interpreter 104, the graphical annotation engine 106, the natural language generation system 108 and/or the user interface 310 may be implemented in software, hardware, firmware, or in some combination to achieve the capabilities described herein.

In the example embodiment shown, computing system 300 comprises a computer memory ("memory") 301, a display 302, one or more processors 303, input/output devices 304 (e.g., keyboard, mouse, CRT or LCD display, touch screen, gesture sensing device and/or the like), other computer-readable media 305, and communications interface 306. The processor 303 may, for example, be embodied as various means including one or more microprocessors with accompanying digital signal processor(s), one or more processor(s) without an accompanying digital signal processor, one or more coprocessors, one or more multi-core processors, one or more controllers, processing circuitry, one or more computers, various other processing elements including integrated circuits such as, for example, an application-specific integrated circuit (ASIC) or field-programmable gate array (FPGA), or some combination thereof. Accordingly, although

illustrated in Figure 3 as a single processor, in some embodiments the processor 303 comprises a plurality of processors. The plurality of processors may be in operative communication with each other and may be collectively configured to perform one or more functionalities of the graphical annotation environment as described herein.

5 The data analyzer 102, the data interpreter 104, the graphical annotation engine 106, the natural language generation system 108 and/or the user interface 310 are shown residing in memory 301. The memory 301 may comprise, for example, transitory and/or non-transitory memory, such as volatile memory, non-volatile memory, or some combination thereof. Although illustrated in Figure 3 as a single memory, the memory 301
10 may comprise a plurality of memories. The plurality of memories may be embodied on a single computing device or may be distributed across a plurality of computing devices collectively configured to function as the graphical annotation environment. In various example embodiments, the memory 301 may comprise, for example, a hard disk, random access memory, cache memory, flash memory, a compact disc read only memory (CD-
15 ROM), digital versatile disc read only memory (DVD-ROM), an optical disc, circuitry configured to store information, or some combination thereof.

 In other embodiments, some portion of the contents, some or all of the components of the data analyzer 102, the data interpreter 104, the graphical annotation engine 106, the natural language generation system 108 and/or the user interface 310
20 may be stored on and/or transmitted over the other computer-readable media 305. The components of the data analyzer 102, the data interpreter 104, the graphical annotation engine 106, the natural language generation system 108 and/or the user interface 310 preferably execute on one or more processors 303 and are configured to generate graphical annotations, as described herein.

25 Alternatively or additionally, other code or programs 330 (e.g., an administrative interface, a Web server, and the like) and potentially other data repositories, such as data repository 340, also reside in the memory 301, and preferably execute on one or more processors 303. Of note, one or more of the components in Figure 3 may not be present in any specific implementation. For example, some embodiments may not provide other
30 computer readable media 305 or a display 302.

 In some example embodiments, as described above, the graphical annotation engine 106 may comprise a scale determination engine 120, an annotation location determiner 122, a graphical output generator 124 or the like. The natural language generation system 108 may comprise a document planner 130, a microplanner 132, a
35 realizer 134 and/or the like. In some example embodiments, the data analyzer 102, the data interpreter 104, the graphical annotation engine 106, the natural language generation system 108 and/or the user interface 310 may also include or otherwise be in

data communication with raw input data 110, event log 112 and/or the domain model 114. The data analyzer 102, the data interpreter 104, the graphical annotation engine 106, the natural language generation system 108 and/or the user interface 310 are further configured to provide functions such as those described with reference to Figure 1.

5 The data analyzer 102, the data interpreter 104, the graphical annotation engine 106, the natural language generation system 108 and/or the user interface 310 may interact with the network 350, via the communications interface 306, with remote data sources 356 (e.g. remote reference data, remote performance data, remote aggregation data and/or the like), third-party content providers 354 and/or client devices 358. The
10 network 350 may be any combination of media (e.g., twisted pair, coaxial, fiber optic, radio frequency), hardware (e.g., routers, switches, repeaters, transceivers), and protocols (e.g., TCP/IP, UDP, Ethernet, Wi-Fi, WiMAX, Bluetooth) that facilitate communication between remotely situated humans and/or devices. In some instance the network 350 may take the form of the internet or may be embodied by a cellular network
15 such as an LTE based network. In this regard, the communications interface 306 may be capable of operating with one or more air interface standards, communication protocols, modulation types, access types, and/or the like. The client devices 358 include desktop computing systems, notebook computers, mobile phones, smart phones, personal digital assistants, tablets and/or the like.

20 In an example embodiment, components/modules of the data analyzer 102, the data interpreter 104, the graphical annotation engine 106, the natural language generation system 108 and/or the user interface 310 are implemented using standard programming techniques. For example, the data analyzer 102, the data interpreter 104, the graphical annotation engine 106, the natural language generation system 108 and/or
25 the user interface 310 may be implemented as a "native" executable running on the processor 303, along with one or more static or dynamic libraries. In other embodiments, the data analyzer 102, the data interpreter 104, the graphical annotation engine 106, the natural language generation system 108 and/or the user interface 310 may be implemented as instructions processed by a virtual machine that executes as one of the
30 other programs 330. In general, a range of programming languages known in the art may be employed for implementing such example embodiments, including representative implementations of various programming language paradigms, including but not limited to, object-oriented (e.g., Java, C++, C#, Visual Basic.NET, Smalltalk, and the like), functional (e.g., ML, Lisp, Scheme, and the like), procedural (e.g., C, Pascal, Ada, Modula, and the like), scripting (e.g., Perl, Ruby, Python, JavaScript, VBScript, and the
35 like), and declarative (e.g., SQL, Prolog, and the like).

The embodiments described above may also use synchronous or asynchronous client-server computing techniques. Also, the various components may be implemented using more monolithic programming techniques, for example, as an executable running on a single processor computer system, or alternatively decomposed using a variety of structuring techniques, including but not limited to, multiprogramming, multithreading, client-server, or peer-to-peer, running on one or more computer systems each having one or more processors. Some embodiments may execute concurrently and asynchronously, and communicate using message passing techniques. Equivalent synchronous embodiments are also supported. Also, other functions could be implemented and/or performed by each component/module, and in different orders, and by different components/modules, yet still achieve the described functions.

In addition, programming interfaces to the data stored as part of the data analyzer 102, the data interpreter 104, the graphical annotation engine 106, the natural language generation system 108 and/or the user interface 310, such as by using one or more application programming interfaces can be made available by mechanisms such as through application programming interfaces (API) (e.g. C, C++, C#, and Java); libraries for accessing files, databases, or other data repositories; through scripting languages such as XML; or through Web servers, FTP servers, or other types of servers providing access to stored data. The raw input data 110, the event log 112 and the domain model 114 may be implemented as one or more database systems, file systems, or any other technique for storing such information, or any combination of the above, including implementations using distributed computing techniques. Alternatively or additionally, the raw input data 110, the event log 112 and the domain model 114 may be local data stores but may also be configured to access data from the remote data sources 356.

Different configurations and locations of programs and data are contemplated for use with techniques described herein. A variety of distributed computing techniques are appropriate for implementing the components of the illustrated embodiments in a distributed manner including but not limited to TCP/IP sockets, RPC, RMI, HTTP, Web Services (XML-RPC, JAX-RPC, SOAP, and the like). Other variations are possible. Also, other functionality could be provided by each component/module, or existing functionality could be distributed amongst the components/modules in different ways, yet still achieve the functions described herein.

Furthermore, in some embodiments, some or all of the components of the data analyzer 102, the data interpreter 104, the graphical annotation engine 106, the natural language generation system 108 and/or the user interface 310 may be implemented or provided in other manners, such as at least partially in firmware and/or hardware, including, but not limited to one or more ASICs, standard integrated circuits, controllers

executing appropriate instructions, and including microcontrollers and/or embedded controllers, FPGAs, complex programmable logic devices ("CPLDs"), and the like. Some or all of the system components and/or data structures may also be stored as contents (e.g., as executable or other machine-readable software instructions or structured data) on a computer-readable medium so as to enable or configure the computer-readable medium and/or one or more associated computing systems or devices to execute or otherwise use or provide the contents to perform at least some of the described techniques. Some or all of the system components and data structures may also be stored as data signals (e.g., by being encoded as part of a carrier wave or included as part of an analog or digital propagated signal) on a variety of computer-readable transmission mediums, which are then transmitted, including across wireless-based and wired/cable-based mediums, and may take a variety of forms (e.g., as part of a single or multiplexed analog signal, or as multiple discrete digital packets or frames). Such computer program products may also take other forms in other embodiments. Accordingly, embodiments of this disclosure may be practiced with other computer system configurations.

Figures 4-6 illustrate example flowcharts of the operations performed by an apparatus, such as computing system 300 of Figure 3, in accordance with example embodiments of the present invention. It will be understood that each block of the flowcharts, and combinations of blocks in the flowcharts, may be implemented by various means, such as hardware, firmware, one or more processors, circuitry and/or other devices associated with execution of software including one or more computer program instructions. For example, one or more of the procedures described above may be embodied by computer program instructions. In this regard, the computer program instructions which embody the procedures described above may be stored by a memory 301 of an apparatus employing an embodiment of the present invention and executed by a processor 303 in the apparatus. As will be appreciated, any such computer program instructions may be loaded onto a computer or other programmable apparatus (e.g., hardware) to produce a machine, such that the resulting computer or other programmable apparatus provides for implementation of the functions specified in the flowcharts' block(s). These computer program instructions may also be stored in a non-transitory computer-readable storage memory that may direct a computer or other programmable apparatus to function in a particular manner, such that the instructions stored in the computer-readable storage memory produce an article of manufacture, the execution of which implements the function specified in the flowcharts' block(s). The computer program instructions may also be loaded onto a computer or other programmable apparatus to cause a series of operations to be performed on the computer or other

programmable apparatus to produce a computer-implemented process such that the instructions which execute on the computer or other programmable apparatus provide operations for implementing the functions specified in the flowcharts' block(s). As such, the operations of Figures 4-6, when executed, convert a computer or processing circuitry
5 into a particular machine configured to perform an example embodiment of the present invention. Accordingly, the operations of Figures 4-6 define an algorithm for configuring a computer or processor, to perform an example embodiment. In some cases, a general purpose computer may be provided with an instance of the processor which performs the algorithm of Figures 4-6 to transform the general purpose computer into a particular
10 machine configured to perform an example embodiment.

Accordingly, blocks of the flowchart support combinations of means for performing the specified functions and combinations of operations for performing the specified functions. It will also be understood that one or more blocks of the flowcharts', and combinations of blocks in the flowchart, can be implemented by special purpose
15 hardware-based computer systems which perform the specified functions, or combinations of special purpose hardware and computer instructions.

In some example embodiments, certain ones of the operations herein may be modified or further amplified as described below. Moreover, in some embodiments additional optional operations may also be included (some examples of which are shown
20 in dashed lines in Figure 4). It should be appreciated that each of the modifications, optional additions or amplifications described herein may be included with the operations herein either alone or in combination with any others among the features described herein.

Figure 4 is a flow chart illustrating an example method for generating graphical
25 annotations. As is shown in operation 402, an apparatus may include means, such as the data analyzer 102, the graphical annotation engine 106, the processor 303, or the like, for receiving an indication of an alarm condition. In some example embodiments an alarm may cause the selection of a primary data channel and a determination of a time period in which the alarm was generated. Alternatively or additionally other means may be used to
30 alert the apparatus to a primary data channel, such as, but not limited to, a user action, a detected pattern in the raw input data or a data channel, a determined value in the raw input data or a data channel, and/or the like.

As is shown in operation 404, an apparatus may include means, such as the data analyzer 102, the data interpreter 104, the processor 303, or the like, for determining one
35 or more key patterns in a primary data channel. In some example embodiments the key patterns may be determined based on the time period of the alarm condition, however in

other examples a larger or smaller time period may be selected. The determination of the one or more key patterns is further described with reference to Figure 5.

As is shown in operation 406, an apparatus may include means, such as the data analyzer 102, the data interpreter 104, the processor 303, or the like, for determining one or more significant patterns in one or more related data channels. In some example embodiments, the apparatus, such as via the data analyzer 102 may determine one or related channels based on one or more predefined relationships. In some examples, the predefined relationships may be defined by the domain model 114. The determination of the one or more significant patterns is further described with reference to Figure 6.

As is shown in operation 408, an apparatus may include means, such as the data analyzer 102, the data interpreter 104, the processor 303, or the like, for determining one or more contextual channels to be included in the graphical output. The one or more contextual channels may provide events or other context that may be indicative of the cause of the one or more key patterns and/or the one or more significant patterns. As is shown in operation 410, an apparatus may include means, such as the graphical annotation engine 106, the scale determination engine 120, the processor 303, or the like, for determining a time period to be represented by the graphical output. In some example embodiments, the time period chosen for the graph is the time period in which the one or more key patterns are displayed. As is shown in operation 412, an apparatus may include means, such as the natural language generation system 108, the document planner 130, the microplanner 132, the realizer 134, the processor 303, or the like, for generating a natural language annotation of at least one of the one or more key patterns or the one or more significant patterns.

As is shown in operation 414, an apparatus may include means, such as the graphical annotation engine 106, the annotation location determiner 122, the graphical output generator 124, the processor 303, the user interface 310 or the like, for generating a graphical output that is configured to be displayed in a user interface. In some example embodiments, the graph is configured to utilize the determined scale to display the primary data channel, one or more related channels having significant events, natural language annotations, a narrative, events and/or the like. In some example embodiments and in an instance in which a user clicks on a text annotation in the graph, a corresponding phrase in the situation analysis text may be highlighted and/or in an instance in which a user clicks on underlined phrase in the narrative or situation analysis text, a corresponding annotation may be highlighted on the graph.

Figure 5 is a flow chart illustrating an example method determining one or more key patterns in a primary data channel. As is shown in operation 502, an apparatus may include means, such as the data analyzer 102, the data interpreter 104, the processor

303, or the like, for identifying one or more patterns wherein a pattern is at least one of a trend, spike or step in the data channel. As is shown in operation 504, an apparatus may include means, such as the data analyzer 102, the data interpreter 104, the processor 303, or the like, for assigning an importance level to the one or more patterns. As is
5 shown in operation 506, an apparatus may include means, such as the data analyzer 102, the data interpreter 104, the processor 303, or the like, for identifying one or more key patterns of the one or more patterns, wherein a key pattern is a pattern that exceeds a predefined importance level.

Figure 6 is a flow chart illustrating an example method determining one or more
10 significant patterns in a related data channel. As is shown in operation 602, an apparatus may include means, such as the data analyzer 102, the data interpreter 104, the processor 303, or the like, for identifying one or more unexpected patterns in a related data channel in response to detecting one or more patterns in the data channel. In some example embodiments, the one or more patterns identified in an another data channel
15 violate a predetermined constraint, threshold or the like may be considered as unexpected patterns or anomalies. The one or patterns may be identified within the time period or scale identified for the graph. As is shown in operation 604, an apparatus may include means, such as the data analyzer 102, the data interpreter 104, the processor 303, or the like, for assigning an importance level to the one or more unexpected
20 patterns. As is shown in operation 606, an apparatus may include means, such as the data analyzer 102, the data interpreter 104, the processor 303, or the like, for identifying one or more significant patterns of the one or more unexpected patterns, wherein a significant pattern is an unexpected pattern.

Many modifications and other embodiments of the inventions set forth herein will
25 come to mind to one skilled in the art to which these inventions pertain having the benefit of the teachings presented in the foregoing descriptions and the associated drawings. Therefore, it is to be understood that the inventions are not to be limited to the specific embodiments disclosed and that modifications and other embodiments are intended to be included within the scope of the appended claims. Moreover, although the foregoing
30 descriptions and the associated drawings describe example embodiments in the context of certain example combinations of elements and/or functions, it should be appreciated that different combinations of elements and/or functions may be provided by alternative embodiments without departing from the scope of the appended claims. In this regard, for example, different combinations of elements and/or functions than those explicitly
35 described above are also contemplated as may be set forth in some of the appended claims. Although specific terms are employed herein, they are used in a generic and descriptive sense only and not for purposes of limitation.

THAT WHICH IS CLAIMED:

1. A method comprising:
 - detecting one or more patterns in a data channel derived from raw input data;
 - identifying one or more patterns in another data channel also derived from the
5 raw input data;
 - generating one or more phrases describing the one or more patterns in the
data channel and the one or more patterns in the another data channel;
 - and
 - generating a graphical output based on the data channel, the another data
10 channel and the one or more phrases, wherein the one or more phrases
are interactively annotated on the graphical output of the data channel and
the another data channel.

2. A method according to Claim 1, further comprising:
 - detecting one or more patterns in the data channel by:
15 identifying one or more patterns wherein a pattern is at least one of
a trend, spike or step in the data channel;
 - assigning an importance level to the one or more patterns; and
 - identifying one or more key patterns of the one or more patterns,
wherein a key pattern is a pattern that exceeds a predefined
20 importance level.

3. A method according to Claim 2, further comprising:
 - determining a time period to be displayed in a graph, wherein the time
period chosen for the graph is the time period in which the one or more
key patterns are displayed.

- 25 4. A method according to Claim 3, further comprising:
 - identifying one or more patterns in the another data channel by:
 - assigning an importance level to one or more unexpected patterns;
 - and
 - 30 identifying one or more significant patterns of the one or more
unexpected patterns, wherein a significant pattern is an
unexpected pattern.

5. A method according to Claim 4, further comprising:

determining that the one or more patterns identified in an another data channel violate a predetermined constraint; and determining that the one or more patterns are one or more unexpected patterns.

- 5 6. A method according to Claim 4, wherein the one or more unexpected patterns are identified within the time period of the graph.
7. A method according to Claim 4, further comprising:
generating the graph based on the graphical output, wherein the graph comprises at least a portion of the data channel that
10 contains the one or more key events, the another data channel that contains the one or more significant events and the one or more contextual channels, the at least a portion of the one or more key events, the one or more significant events and the one or more contextual channels being annotated by the one or
15 more phrases;
generating a narrative; and
causing the graph to be displayed in a user interface.
8. A method according to Claim 7, further comprising:
determining the one or more contextual channels to be displayed on the
20 graph, wherein the one or more contextual channels are configured to provide context information for at least one of the one or more patterns in the data channel or the one or more patterns in the another data channel.
9. A method according to Claim 1, wherein the data channel is selected from one
25 or more data channels derived from the raw input data, the data channel is selected based on an indication received from at least one of a user, an alarm or another computing device.
10. A method according to Claim 7, wherein additional information related to the one or more phrases is displayed in response to a user interaction with the user interface.
- 30 11. A method according to Claim 7, wherein the one or more key patterns and the one or more significant patterns are highlighted in response to a user interaction with the one or more phrases in the user interface.

12. An apparatus comprising:

at least one processor; and

at least one memory including computer program code, the at least one memory and the computer program code configured to, with the at least one processor, cause the apparatus to at least:

determine one or more key patterns in a primary data channel, wherein the primary data channel is derived from raw input data in response to a constraint being satisfied;

determine one or more significant patterns in one or more related data channels;

generate a natural language annotation for at least one of the one or more key patterns or the one or more significant patterns;

and

generate a graph that is configured to be displayed in a user interface, the graph visually presenting at least a portion of the primary data channel having one or more key patterns, at least a portion of the another data channel having the one or more significant patterns and the natural language annotation.

13. An apparatus according to Claim 12, wherein the at least one memory including the computer program code is further configured to, with the at least one processor, cause the apparatus to:

identify one or more patterns in the primary data channel;

assign an importance level to the one or more patterns; and

identify one or more key patterns of the one or more patterns in an

instance in which the importance level of a pattern of the one or more patterns exceeds a threshold defined by a domain model.

14. An apparatus according to Claim 13, wherein the at least one memory including the computer program code is further configured to, with the at least one processor, cause the apparatus to:

determine a time period to be represented by the graphical output, wherein the time period is configured such that at least a portion of the one or more key patterns occur within the time period.

15. An apparatus according to Claim 13, wherein the at least one memory including the computer program code is further configured to, with the at least one processor, cause the apparatus to:

determine one or more patterns in one or more related data channels that corresponds to an occurrence of the one or more key patterns; and determine one or more significant patterns of the one or more patterns, wherein a significant pattern is a pattern that represents an anomaly.

5 16. An apparatus according to Claim 15, wherein an anomaly is defined by the domain model.

10 17. An apparatus according to Claim 12, wherein the natural language annotation is generated by a natural language generation system based on the raw input data and at least a portion of the one or more key patterns and the one or more significant patterns.

18. An apparatus according to Claim 12, wherein the at least one memory including the computer program code is further configured to, with the at least one processor, cause the apparatus to:

15 determine one or more contextual channels to be included in the graph, wherein the one or more contextual channels are configured to provide context information for at least one of the one or more key patterns or the one or more significant patterns.

20 19. An apparatus according to Claim 12, wherein additional information related to the natural language annotation is displayed in the user interface in response to a user interaction with the user interface.

20. An apparatus according to Claim 12, wherein the one or more key patterns and the one or more significant patterns in the graph are highlighted in response to a user interaction with the natural language annotation in the user interface.

25 21. A computer program product comprising:

at least one computer readable non-transitory memory medium having program code instructions stored thereon, the program code instructions which when executed by an apparatus causes the apparatus at least to:

30 identify one or more key patterns in a primary data channel, wherein the one or more key patterns have an importance level that exceeds a predefined importance level; generate a natural language annotation of at least a portion of the one or more key patterns; and

generate a graph for display in a user interface, the graph having at least a portion of the primary data channel having one or more key patterns and the natural language annotation.

22. A computer program product according to Claim 21, further comprising program code instructions configured to:

identify one or more patterns in the primary data channel; and
assign the importance level to the one or more patterns.

23. A computer program product according to Claim 21, further comprising program code instructions configured to:

10 determine one or more patterns in one or more related data channels that occur within the same time period as the one or more key patterns; assign an importance level to the one or more patterns; and determine one or more significant patterns of the one or more patterns.

24. A computer program product according to Claim 23, further comprising program code instructions configured to:

15 determine a time period for the graph, wherein the time period is configured such that at least a portion of the one or more key patterns occur within the time period.

25. A computer program product according to Claim 24, wherein a significant pattern of the one or more significant patterns represents an anomaly in an instance in which the significant pattern is compared with a key pattern of the one or more key patterns that occurs within the same time period.

26. A computer program product according to Claim 25, further comprising program code instructions configured to:

25 generate a natural language annotation of the one or more significant patterns; and
cause at least a portion of the one or more related data channels having the one or more significant patterns and the natural language annotation to be displayed via the graph.

30 27. A computer program product according to Claim 21, wherein the natural language annotation is generated by a natural language generation system based on raw input data corresponding to at least a portion of the one or more key patterns and the one or more significant patterns.

28. A computer program product according to Claim 21, further comprising program code instructions configured to:

5 determine one or more contextual channels to be included in the graph,
 wherein the one or more contextual channels are configured to provide
 context information for at least one of the one or more key patterns or
 the one or more significant patterns.

29. A computer program product according to Claim 21, wherein the primary data channel is selected from one or more data channels derived from the raw input data based on an alarm condition.

10 30. A computer program product according to Claim 21, wherein additional information related to the natural language annotation is displayed in response to a user interaction with the graph in the user interface.

15 31. A computer program product according to Claim 21, wherein the one or more key patterns and the one or more significant patterns are highlighted in response to a user interaction with the natural language annotation in the user interface.

32. A computer program product according to Claim 21, further comprising program code instructions configured to:

20 generate one or more messages that correspond to at least a portion of
 the one or more key patterns;
 generate a tree-structured object that contains at least a portion of the one
 or more messages;
 refine the tree-structured object; and
 generate the natural language annotation based on a realization of the
 tree-structured object.

25

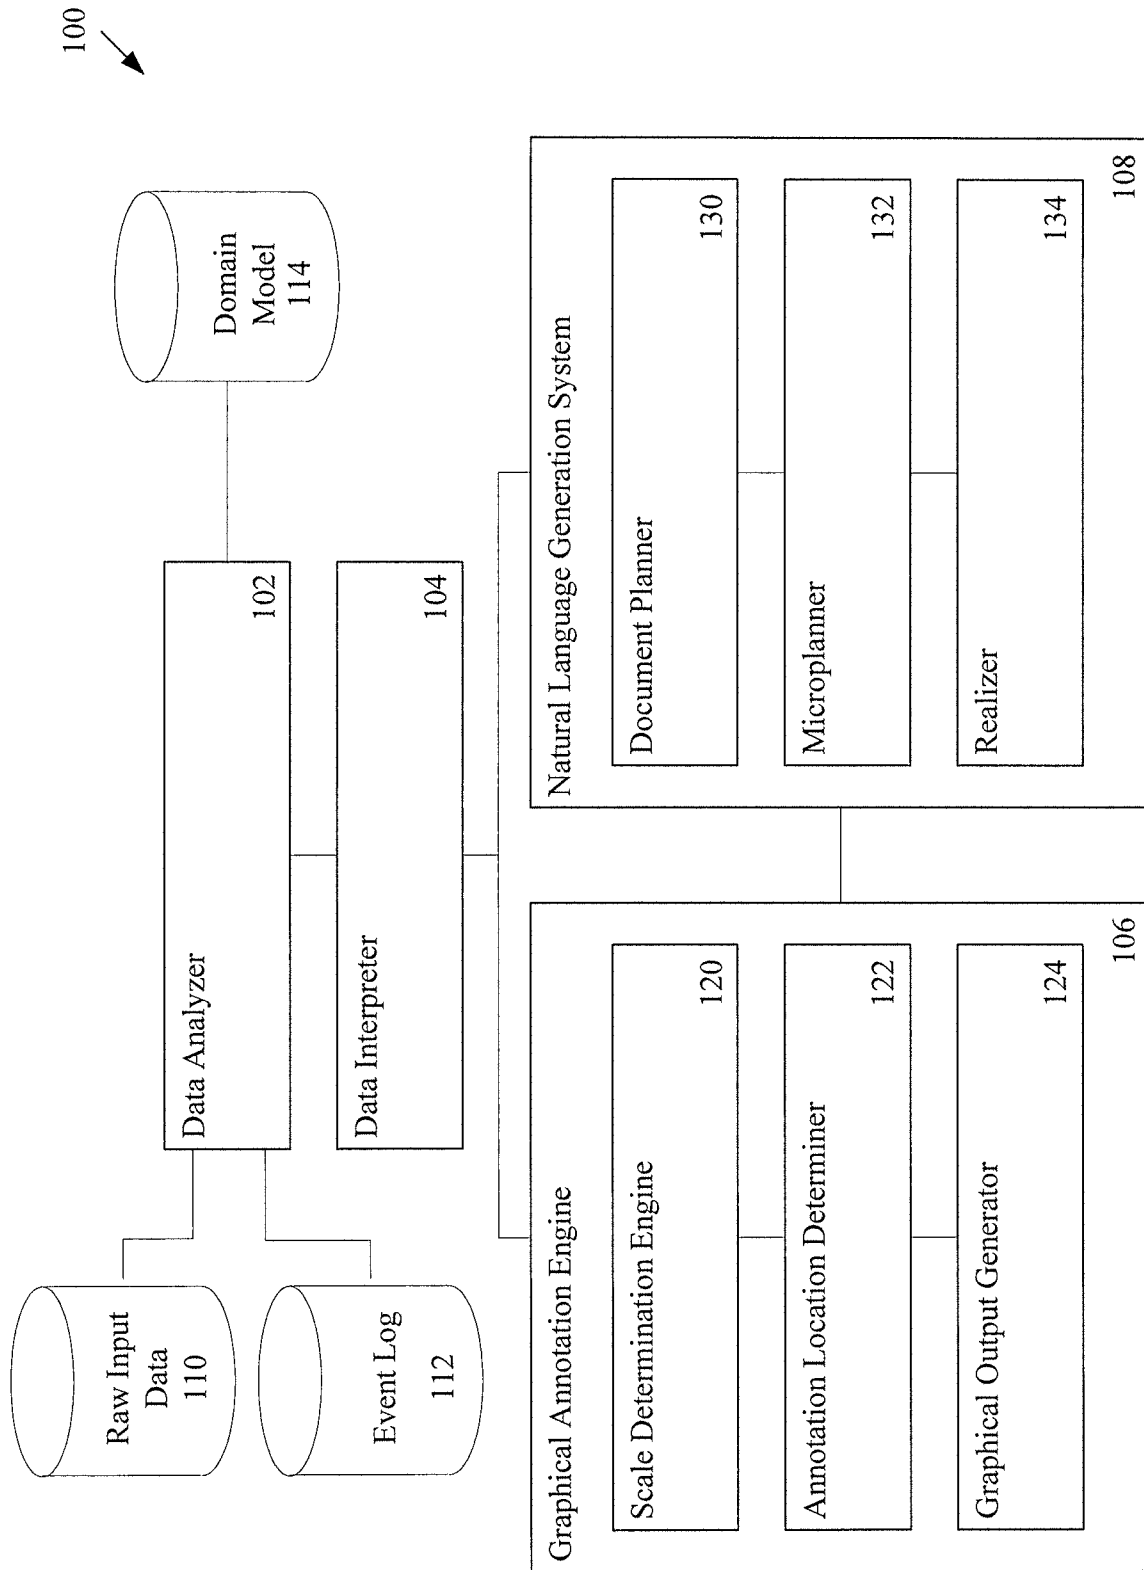


FIG. 1

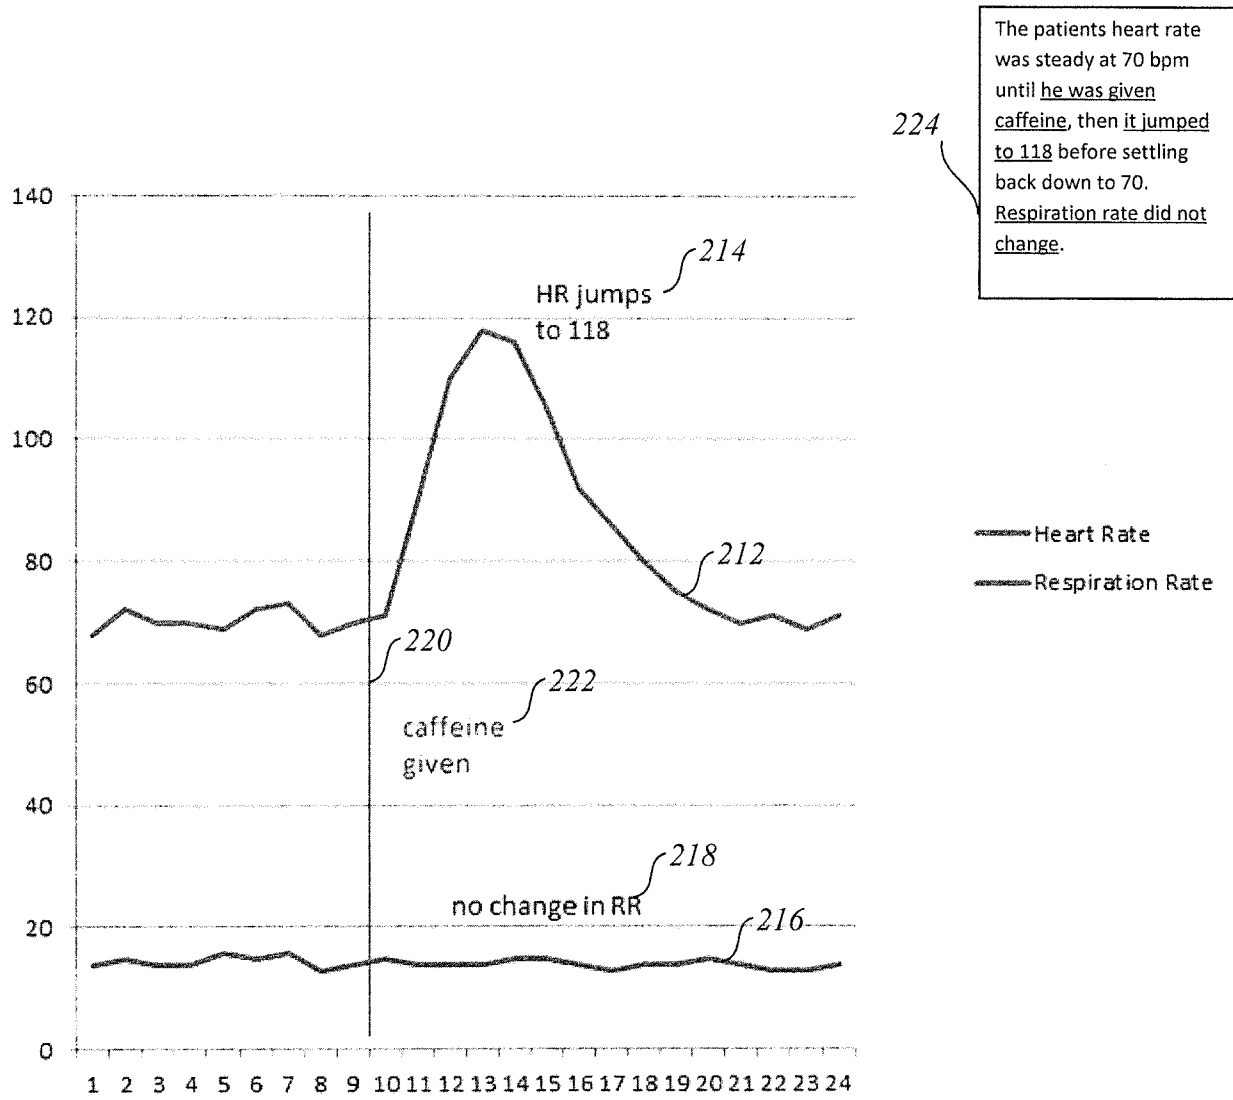


FIG. 2

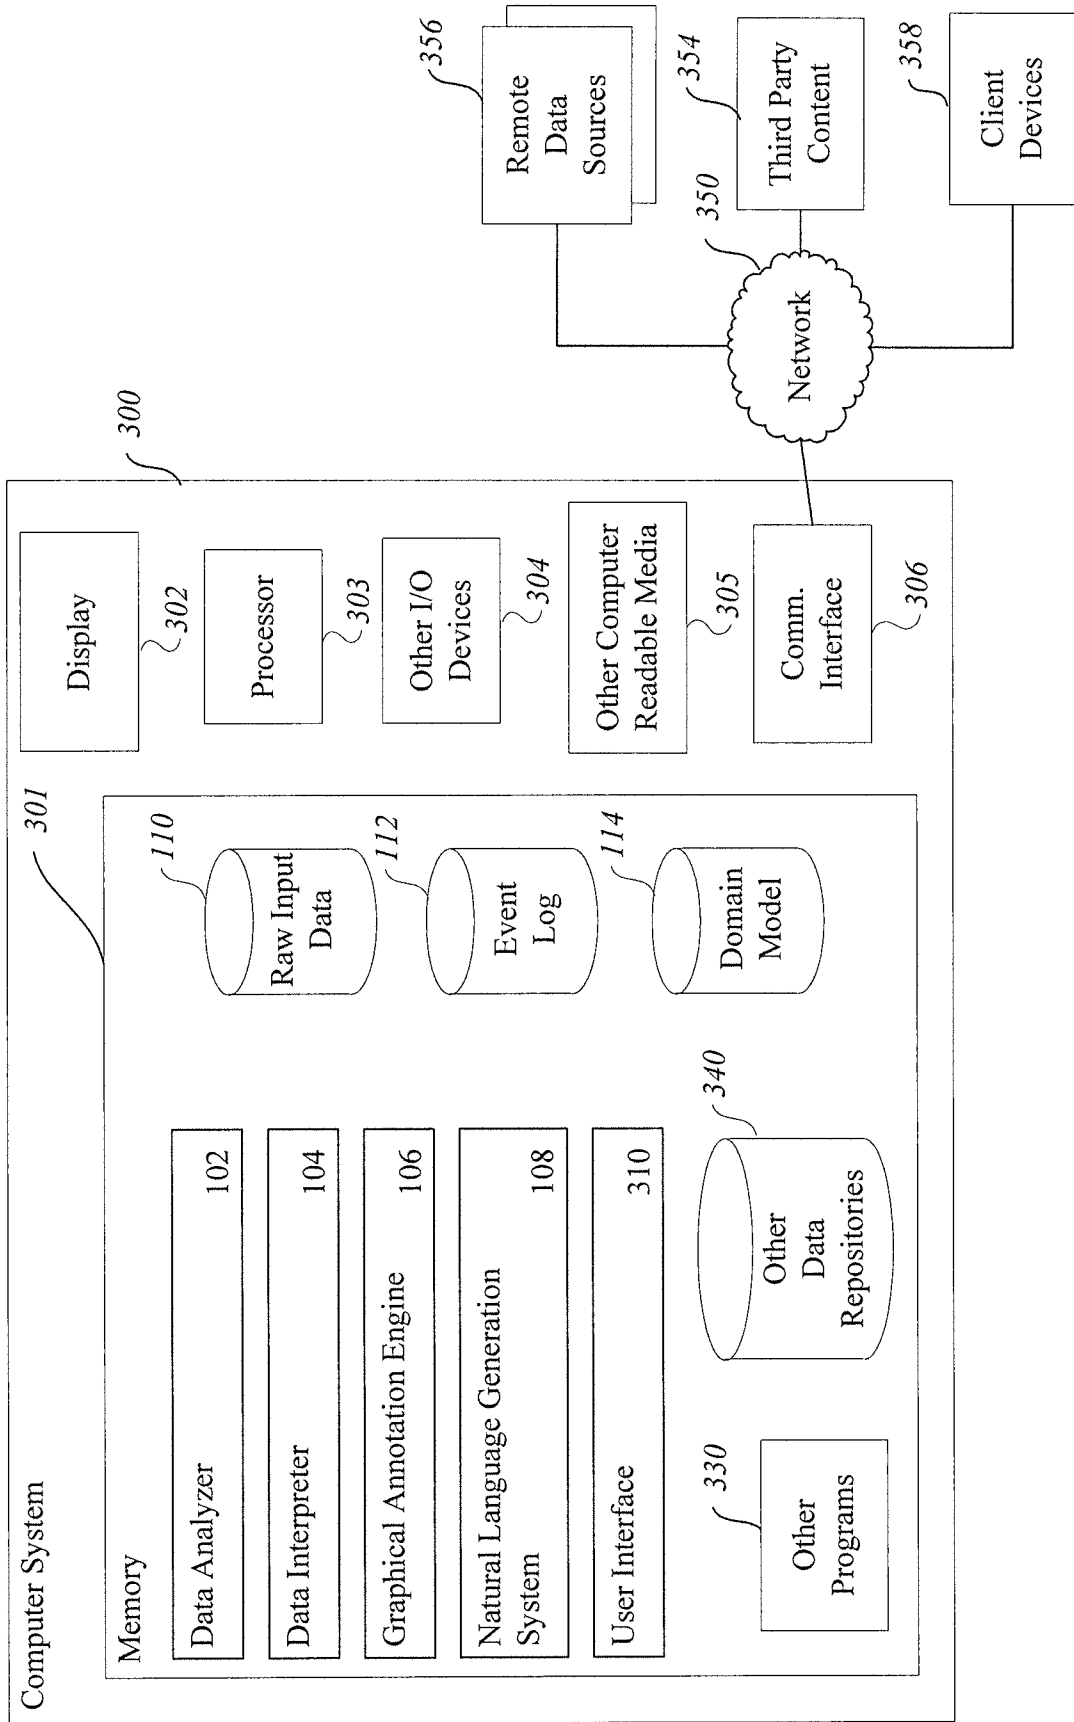
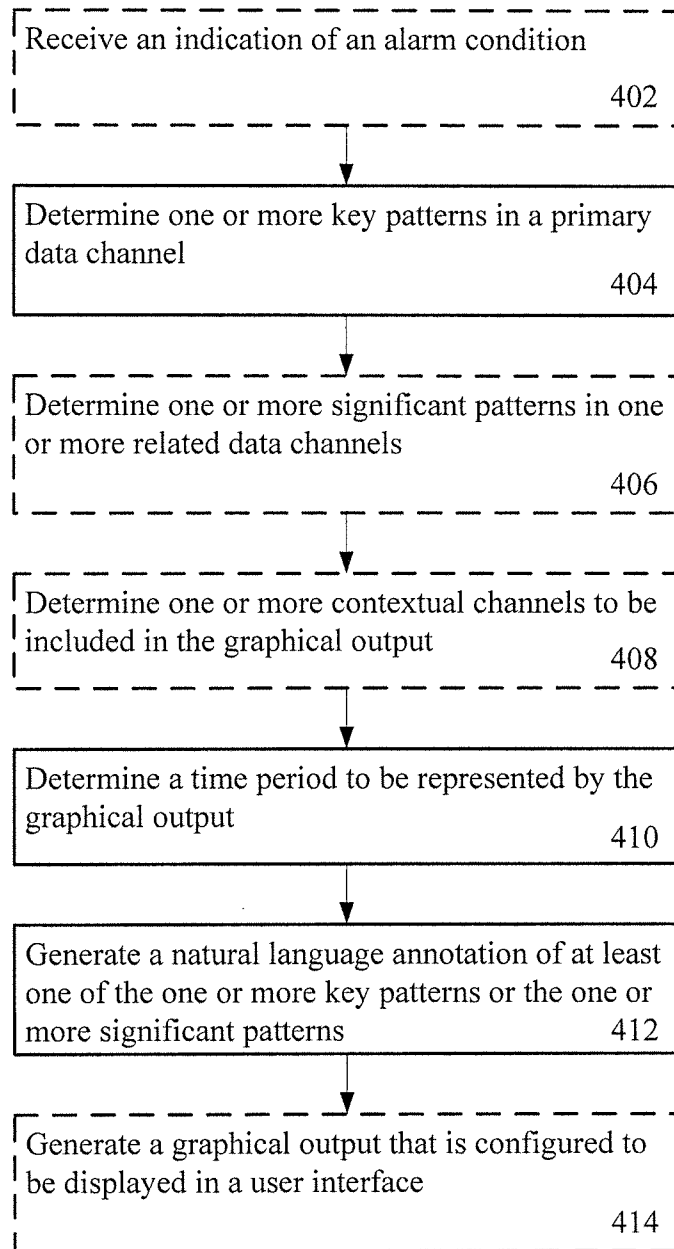
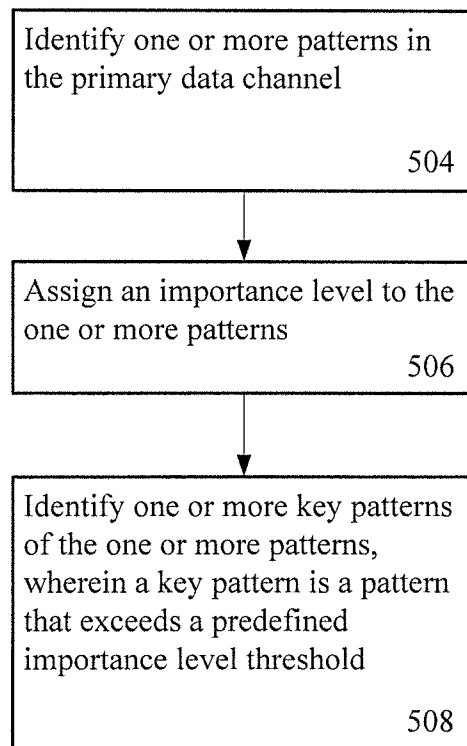
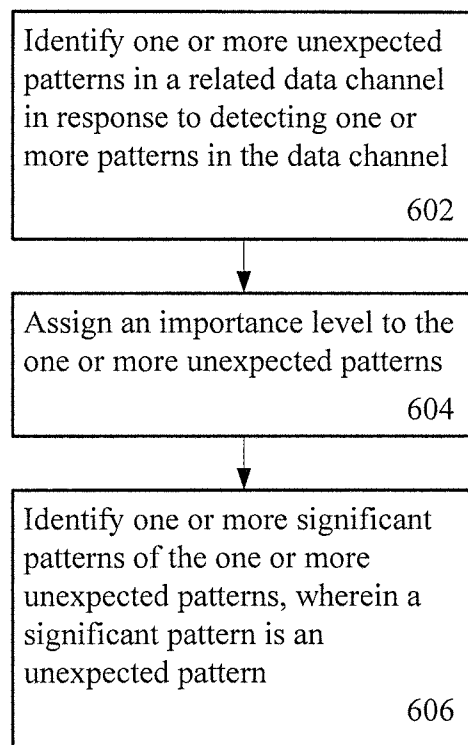


FIG. 3

**FIG. 4**

**FIG. 5**

6/6

**FIG. 6**

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2012/053128

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F17/24
ADD.
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED
Minimum documentation searched (classification system followed by classification symbols)
G06F
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2003/131315 A1 (ESCHER RICHARD E A [CA]) 10 July 2003 (2003-07-10) abstract figures 2,3 paragraph [0024] paragraph [0033] - paragraph [0035] paragraph [0046]	1-32
A	US 2007/129942 A1 (BAN OLIVER K [US] ET AL) 7 June 2007 (2007-06-07) abstract figures 3-5 paragraph [0006] - paragraph [0011] ----- -/--	1,12,21

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E" earlier application or patent but published on or after the international filing date	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O" document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search 20 June 2013	Date of mailing of the international search report 27/06/2013
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Triest, Johannes

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2012/053128

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>US 2005/256703 A1 (MARKEL STEVEN O [US]) 17 November 2005 (2005-11-17) abstract figures 16,17 paragraph [0009] - paragraph [0010] paragraph [0061] - paragraph [0062] -----</p>	1-32
A	<p>WO 2006/010044 A2 (GAPYX LLC [US]; MCNAIR DOUGLAS [US]) 26 January 2006 (2006-01-26) abstract figures 3,4 paragraph [0133] -----</p>	1,12,21

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2012/053128

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2003131315 A1	10-07-2003	CA 2414620 A1 US 2003131315 A1	17-06-2003 10-07-2003

US 2007129942 A1	07-06-2007	NONE	

US 2005256703 A1	17-11-2005	NONE	

WO 2006010044 A2	26-01-2006	US 2006036542 A1 WO 2006010044 A2	16-02-2006 26-01-2006
