



- (51) **International Patent Classification:**
G06T 7/20 (2006.01)
- (21) **International Application Number:**
PCT/IB2012/057774
- (22) **International Filing Date:**
27 December 2012 (27.12.2012)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (71) **Applicant:** **ARRIA DATA2TEXT LIMITED** [GB/GB];
Meston Building, University of Aberdeen, Aberdeen AB24
3UE (GB).
- (72) **Inventor:** **SRIPADA, Gowri, Somayajulu**; 9 Brodiach
Court, Westhill AB32 6QY (GB).
- (74) **Agents:** **CARTER, Stephen** et al; Mewburn Ellis LLP, 33
Gutter Lane, London EC2V 8AS (GB).
- (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, PA, 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).

Declarations under Rule 4.17:

— of inventorship (Rule 4.17(iv))

Published:

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

(54) **Title:** METHOD AND APPARATUS FOR MOTION DESCRIPTION

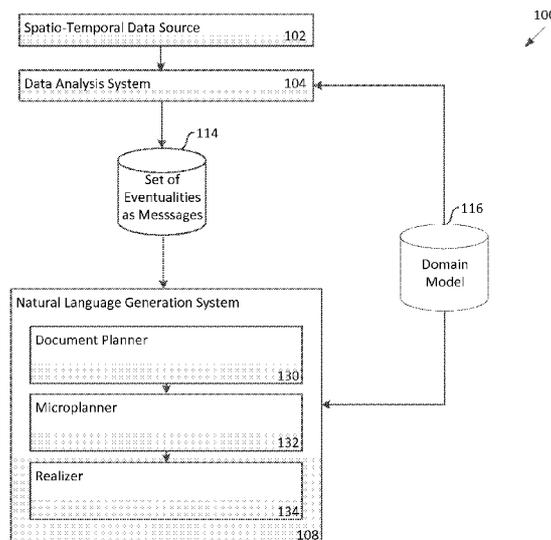


FIG. 1

(57) **Abstract:** A method, apparatus, and computer program product for describing motion. The method may include receiving a set of eventualities (114). The set of eventualities (114) may describe at least one of a domain event and a domain state. The at least one of the domain event and the domain state may be derived from a set of spatio-temporal data (102) and the set of eventualities (114) may be associated with a particular region and a particular time period. The method may include organizing the set of eventualities to generate a document plan. The method may further include generating, using a processor, a linguistic representation of the set of eventualities using the document plan.

WO 2014/102569 A1

METHOD AND APPARATUS FOR MOTION DESCRIPTION

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 motion description.

5

BACKGROUND

Advances in computer processor speeds and other performance characteristics have occurred at a rapid pace in recent history, to the point that many human behaviors can now be thoroughly mimicked by machines. However, it has become apparent that current technology is insufficient for replication of certain activities. For example, the human brain tends to be quite adept at extracting data and drawing inferences and conclusions from complex sets of data. These inferences and conclusions may be used to describe the data in a way that allows another human to easily understand important events that occur in the data set. One such task that employs these reasoning faculties is the use of language to describe events in a concise, natural manner.

10

In an effort to enable computers and other machines to communicate data in a similar manner to human beings, example embodiments of the invention relate to Natural Language Generation (NLG) systems. These NLG systems function to parse data sets and to identify features within the dataset for communication to users, customers, other computer systems, or the like by expressing the features in a linguistic format. In some examples, a 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."

15

20

25

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 among 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.

In some examples, an NLG system may be configured to linguistically express a certain type of data. For example, the NLG system may be configured to describe sports statistics, financial data, weather data, or the like using terminology and linguistic expressions appropriate for the data set. Different terminology, phraseology, idioms, and the like may be used to describe different types of phenomena, and different data domains may require different analysis techniques for efficient generation of linguistic output. For example, an analysis operation for a set of sports data to generate a game recap may require different data analysis techniques than analysis of weather data to generate a weather forecast. One example of such input data may be presented as a set of eventualities associated with a set of static and moving objects. Current implementations do not provide a proper framework for using such eventualities to express the eventualities in a linguistic format.

25

BRIEF SUMMARY

Some example embodiments of an NLG system may relate to description of motion given a set of eventualities which describe objects and their motion in a given spatial region. For example, a NLG system may be employed to generate a weather forecast when provided with a set of eventualities that describe weather objects, such as areas of precipitation, wind direction, high and low pressure, and the like. These data features may be utilized by the NLG system in conjunction with weather forecast document models to generate the weather forecast in a natural, efficient, linguistically correct format suitable for user consumption.

Methods, apparatuses, and computer program products are described herein that are configured to describe motion. Embodiments may include a method for describing motion. The method may include receiving a set of eventualities. The set of eventualities

may describe at least one of a domain event and a domain state. The at least one of the domain event and the domain state may be derived from a set of spatio-temporal data and the set of eventualities may be associated with a particular region and a particular time period. The method may further include organizing the set of eventualities to generate a document plan, and generating, using a processor, a linguistic representation of the set of eventualities using the document plan.

Embodiments may also include an apparatus for describing motion. The apparatus may include a memory coupled to at least one processor. The processor may be configured to receive a set of eventualities. The set of eventualities may describe at least one of a domain event and a domain state. The at least one of the domain event and the domain state may be derived from a set of spatio-temporal data and the set of eventualities may be associated with a particular region and a particular time period. The processor may be further configured to organize the set of eventualities to generate a document plan, and to generate a linguistic representation of the set of eventualities using the document plan.

Yet further embodiments may include a non-transitory computer readable storage medium comprising instructions for describing motion. The instructions, when executed by a processor, may configure the processor to receive a set of eventualities. The set of eventualities may describe at least one of a domain event and a domain state. The at least one of the domain event and the domain state may be derived from a set of spatio-temporal data and the set of eventualities may be associated with a particular region and a particular time period. The instructions may further configure the processor to organize the set of eventualities to generate a document plan, and to generate a linguistic representation of the set of eventualities using the document plan.

25

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 an motion description system that may benefit from some example embodiments of the present invention;

Figure 2 illustrates a block diagram of an apparatus that embodies a motion description system in accordance with some example embodiments of the present invention;

Figure 3 depicts an illustration of a particular spatial region and corresponding eventualities in accordance with some example embodiments of the present invention; and

Figures 4-6 are flowchart depictions of example methods that may be performed by a motion description system in accordance with some example embodiments of the present invention.

5

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 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.

One of the primary factors that users generally consider in the analysis of spatio-temporal data is the concept of motion. The relative position of attributes of interest over time can help a user determine where the attributes of interest has been, where it is going, how fast it will get there, and other relevant data. From biology to chemistry to physics to meteorology, the movement of attributes of interest within a system can provide valuable information about the system and the attributes of interest within it.

Systems have been developed to detect motion from such sets of spatio-temporal data, and represent moving and static objects as a set of eventualities. The term "eventualities" in the present context is understood to relate to both motion events (e.g., an object moving in a particular direction at a particular speed) and non-motion states (e.g., an object that does not change in position over time). These eventualities may be provided to a system for the purpose of generating a linguistic representation of the data contained within the eventualities. For example, the eventualities may be derived from a set of spatio-temporal data as described in PCT Patent Application No. XX/XXX,XXX, <(Motion Detection Patent Application^ which is incorporated by reference in its entirety herein.

Moving objects and static objects may be characterized as "domain events" and "domain states", respectively by attaching domain-specific cluster motion types (e.g., "spreading", "receding", or "a band of precipitation" for a weather domain, "gridlock", "stop-and-go", or "congested" for a traffic domain, or the like) to the respective object.

Cluster motion types may be assigned to moving and static objects based on the cluster motion types of the motion vectors associated with the objects. Where the constituent motion vectors do not correspond to a simple domain event, the domain event may be classified as a "hybrid movement." In order to express computed domain events and states linguistically by an NLG system, the domain specific cluster motion types may be analyzed to ensure that the domain specific cluster motion types fit into a language friendly ontology of domain events and states. As such, the computed domain events and states may be identified as linguistically describable using words and phrases from the sublanguage used in a specific domain (e.g. a sublanguage for weather reports).

This process is particularly relevant to the field of natural language generation, as other techniques for identifying the motion of objects are not concerned with linguistic expressions of said motion. For example, a robot may be fitted with a computer vision module to drive a vehicle in real-world traffic. Such a robot might compute motion events and states which may not be describable in language, as driving the vehicle is the objective of this robot, but describing the other moving vehicles in linguistic terms is not.

The set of eventualities may be associated with a particular domain. For example, types of motion and terminology used to describe said motion may differ based on the type of data and the associated subject matter. For example, terminology used to describe weather data may differ from terminology used to describe automobile traffic.

Upon analysis of the set of eventualities, embodiments of the invention may function to generate a set of text describing the motion of objects as described herein. These eventualities may be processed to develop a document plan that may be subsequently realized as a fluent and coherent text. For example, precipitation data associated with a geographic region may be analyzed to identify the movement of precipitation systems, and a weather forecast may be generated based on the movement information. Although example embodiments are described with respect to meteorological applications, the systems, apparatuses, methods, and computer products described herein could be equally applicable to analysis and text generation for any set of visual data.

Figure 1 is an example block diagram of example components of an example motion description system 100. In some example embodiments, the motion description system 100 may comprise a data analysis system 104, and a natural language generation system 108. The data analysis system 104, and/or the natural language generation system 108 may take the form of, for example, a code module, a component, circuitry and/or the like. The components of the motion description system 100 are configured to provide various logic (e.g. code, instructions, functions, routines and/or the like) and/or services related to the motion description system.

In some example embodiments a spatio-temporal data source 102 is configured to provide a set of spatio-temporal data, such as from one or more sensors. For example, the spatio-temporal data source 102 may receive information that describes the position of one or more attributes of interest over time. The spatio-temporal data source 102 is operable to provide data that may be used by the data analysis system 104 to generate a set of eventualities represented as messages 114. Example types of spatio-temporal data provided by the spatio-temporal data source 102 may include weather data (e.g., weather sensors), traffic data (e.g., areas of automobile congestion on a street map), scientific data (e.g., growth of cells in a petri dish), medical data (e.g., analysis of an electrocardiograph wave form), or network data (e.g., a visual representation of bandwidth in a network).

The spatio-temporal data source 102 may in some example embodiments be in data communication with one or more sensors, monitoring systems, storage devices, computing nodes, and/or the like. In examples in which the spatio-temporal data source 102 is in communication with a monitoring system or a sensor, the spatio-temporal data source 102 may receive a set of spatial data provided in a visual format that includes one or more images. However, in some example embodiments, the methods, apparatus and computer products described herein may operate without a spatio-temporal data source 102.

The data analysis system 104 may be configured to receive input of raw data, such as raw data received via the spatio-temporal data source 102 and/or other data determined or otherwise obtained by the spatio-temporal data source 102 or the data analysis system 104. Alternatively or additionally, the data analysis system 104 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) in order to detect or otherwise determine the presence of a motion event or a change in spatial data. In other words, in some example embodiments, the data analysis system 104 may function as or otherwise embody the spatio-temporal data source 102.

Spatio-temporal data may include data such as, but not limited to, data that indicates variation across location (e.g. rainfall in different regions), or any data that combines both time series data and spatial data (e.g. rainfall across time in different geographical output areas). The data contained or otherwise made accessible by the spatio-temporal data source 102 may be provided in the form of numeric values (e.g., coordinate 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 data 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 the spatio-temporal data source 102.

5 The data analysis system 104 may identify motion of one or more objects in the set of spatial data provided by the spatio-temporal data source 102 to generate messages that summarize patterns of motion identified in the set of spatio-temporal data. These messages represent a set of eventualities represented as messages 114 that define a series of motion events and/or non-motion states which relate to the motion detected in the set of spatio-temporal data. As such, the output of the data analysis
10 system 104 may be a set of eventualities represented as messages 114. However, it should be appreciated that the data analysis system 104 may also output a set of eventuality data in a format other than messages, and a later process may convert the set of eventuality data into a message format. An example representation of a set of eventualities is described below with respect to Figure 3.

15 The set of eventualities represented as messages 114 may be generated for a particular domain, as defined by a domain model 116. As described above, the domain model 116 may be related to the type of data and/or the data desired to be derived from the spatio-temporal data (e.g., weather data, traffic data, network data, or biological data). The domain model 116 may include terms and data types related to detection and
20 description of moving and static objects within the spatio-temporal data. The domain model 116 may include representations of several pieces of domain specific information or knowledge such as (but not limited to) meta data (e.g. expected ranges of data values, data types, thresholds for categorization of data), ontology (e.g. domain-specific motion types, categories for direction of motion, concepts, their relationships and properties,
25 language friendly messages representing eventualities, relationships and inference rules that operate on messages) and a priori specified importance values of messages. For example, the domain model 116 may include a series of definitions and rules for assigning domain-relevant types to particular instances of detected moving and static objects. The domain model 116 may also include message definitions, rules for
30 associating particular eventualities with particular message types, and rules for determining the relative importance of different eventualities. The domain model 116 may contain information related to a particular domain or industry. In some examples, the domain model 116 may provide importance levels, information related to particular behaviors and/or the like. In other examples, the domain model 116 may describe
35 relationships between various events and/or phenomena. For example in a weather domain, a domain model may indicate or otherwise instantiate an extreme weather message in an instance in which 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 are present in the data. The extreme weather message may then be labeled as important, whereas typical temperatures or a typical wind message may not be marked as important in some examples.

5 The set of eventualities represented as messages 114 may be used by the natural language generation system 108 to organize these messages into a document plan. 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
10 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 observed event) that could be expressed via a simple sentence (although it may ultimately be realized by some other linguistic means such as a phrase or word). For example, to linguistically describe
15 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 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
20 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 undertaken as is described herein. An example method for generation of natural language from a set of eventualities using a natural language generation system 108 is
25 described further below with respect to Figure 4.

 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 text. The natural language generation system 108 comprises a document planner 130, a microplanner 132 and/or a
30 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 University Press (2000), which is incorporated by reference in its entirety herein.

 The document planner 130 may be configured to create, filter, organize, and order
35 one or more messages based on the set of set of eventualities represented as messages 114. The document planner 130 may be further configured to determine how to arrange those messages to describe the set of eventualities. The document planner 130 may

comprise a content determination process that is configured to select the messages to provide a complete picture of the provided data. Returning to the weather example, a forecast may be generated that describes all weather features in a particular geographic region or according to a particular day of the week. Various other message orders may be employed based on the presence of absence of particular events. Example methods of generating a document plan from a set of eventualities are described below with respect to Figures 4-6.

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 (e.g. elaboration, consequence, contrast and/or the like) to each other.

The microplanner 132 may be 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 events may be aggregated so that both of the events can be described by a single sentence.

In some examples, lexicalization includes, but is not limited to, choosing particular words for the expression of concepts and relations. For example, the phrase "it will become mainly cloudy across the west" may be used to describe cloud cover moving into the western sky.

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. For example, in a first sentence "rain" and a location may be referenced where "the showers" or "the overcast weather" may be referenced in subsequent sentences.

The output of the microplanner 132, in some example embodiments, is a tree-structured realization specification whose leaf-nodes are ontological properties derived from the data analysis system 104, and whose internal nodes express rhetorical relations between the leaf nodes. A phrase specification may correspond to a sentence or a sub-sentence fragment (e.g. a title) and are produced from one or more messages. A phrase specification is configured to contain one or more syntactic constituents (e.g. subject, verb, prepositional phrase and/or the like) and one or more syntactic features (e.g. tense).

A realizer 134 may be configured to traverse a text specification output by the microplanner 132 to express the text specification in natural language. The realization process that is applied to each phrase specification in the text specification makes use of a grammar (e.g. the grammar of the linguistic resources 114) which specifies the valid syntactic constituents in the language and further provides a way of mapping from phrase specifications 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 output of the realizer 134, in some example embodiments, is a description of the objects and the motion of the objects, or other data that may be derived from the motion information used to build the ontological properties. The realizer may also output text or a narrative that is configured to describe or otherwise summarize the motion of the objects, and/or one or more events associated with the motion of the objects.

Figure 2 is an example block diagram of an example computing device for practicing embodiments of an example motion description system. In particular, Figure 2 shows a computing system 200 that may be utilized to implement a motion description system 100 having a spatio-temporal data source 102; a data analysis system 104, a natural language generation system 108 including, in some examples, a document planner 130, a microplanner 132 and/or a realizer 134, and/or an optional user interface (not shown). One or more general purpose or special purpose computing systems/devices may be used to implement the data analysis system 104, and/or the natural language generation system 108. In addition, the computing system 200 may comprise one or more distinct computing systems/devices and may span distributed locations. In some example embodiments, the data analysis system 104, and/or the natural language generation system 108 may be configured to operate remotely via the network 250. 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 analysis system 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 analysis system 104, and/or the natural language generation system 108 may be implemented in software, hardware, firmware, or in some combination to achieve the capabilities described herein.

In the example embodiment shown, computing system 200 comprises a computer memory ("memory") 201, a display 202, one or more processors 203, input/output devices 204 (e.g., keyboard, mouse, CRT or LCD display, touch screen, gesture sensing

device and/or the like), other computer-readable media 205, and communications interface 206. The processor 203 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 2 as a single processor, in some embodiments the processor 203 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 motion description system as described herein.

The data analysis system 104, and/or the natural language generation system 108 are shown residing in memory 201. The memory 201 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 2 as a single memory, the memory 201 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 motion description system. In various example embodiments, the memory 201 may comprise, for example, a hard disk, random access memory, cache memory, flash memory, a compact disc read only memory (CD-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 analysis system 104, and/or the natural language generation system 108 may be stored on and/or transmitted over the other computer-readable media 205. The components of the data analysis system 104, and/or the natural language generation system 108 preferably execute on one or more processors 203 and are configured to generate natural language describing motion vectors derived from spatial data, as described herein.

Alternatively or additionally, other code or programs 230 (e.g., an administrative interface, a Web server, and the like) and potentially other data repositories also reside in the memory 201, and preferably execute on one or more processors 203. Of note, one or more of the components in Figure 2 may not be present in any specific implementation. For example, some embodiments may not provide other computer readable media 205 or a display 202.

The data analysis system 104, and/or the natural language generation system 108 are further configured to provide functions such as those described with reference to Figure 1. The data analysis system 104, and/or the natural language generation system 108 may interact with the network 250, via the communications interface 206, with the spatio-temporal data source(s) 102, (e.g. remote reference data, remote performance data, remote aggregation data, remote sensors and/or the like), third-party content providers 254 and/or client devices 258. The network 250 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 250 may take the form of the internet or may be embodied by a cellular network such as an LTE based network. In this regard, the communications interface 206 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 258 include desktop computing systems, notebook computers, mobile phones, smart phones, personal digital assistants, tablets and/or the like.

In an example embodiment, components/modules of the data analysis system 104, and/or the natural language generation system 108 are implemented using standard programming techniques. For example, the data analysis system 104 and/or the natural language generation system 108 may be implemented as a "native" executable running on the processor 203, along with one or more static or dynamic libraries. In other embodiments, the data analysis system 104, and/or the natural language generation system 108 may be implemented as instructions processed by a virtual machine that executes as one of the other programs 230. 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 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 analysis system 104, and/or the natural language generation system 108, 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 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.

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 spatio-temporal data source 102, the data analysis system 104, and/or the natural language generation system 108 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
 5 embodiments. Accordingly, embodiments of this disclosure may be practiced with other computer system configurations.

Figure 3 depicts an illustration of a particular spatial region 300 and corresponding eventualities A 302, B 304, C 306, D 308, and E 310 in accordance with some example
 10 embodiments of the present invention. The spatial region 300 and corresponding eventualities 302-310 provide a visual representation of data contained within a set of eventualities, such as the set of set of eventualities represented as messages 114 described with respect to Figures 1 and 2. For example, although the eventualities are depicted in a visual manner with respect to Figure 3, these eventualities may represent a data output from a process for detecting motion in a set of spatio-temporal data. As a
 15 specific example, motion events (e.g., a moving object) and motion states (e.g., a static object) may be depicted as entries in a table, as data vectors, or in any other format which describes motion as derived from a set of spatio-temporal data.

The region 300 depicts a series of weather events throughout the region. As described above, although weather systems are used as an example, the same and
 20 similar methods as described herein may be applied to other domains. In the instant example, eventuality A 302 represents a hurricane system moving in a westerly direction across the region. Eventuality B 304 represents a small weather system (e.g., rain, snow, sleet, or other precipitation) traveling in a southerly direction. Eventuality C 306 represents a spreading weather system traveling in a southeasterly direction. Eventuality
 25 D 308 represents a large weather system traveling in a southerly direction. Eventuality E 310 represents a stationary weather system.

As an example, a data or message structure defining the eventualities 302-310 might be provided according to the following format:

Identifier	Direction	Speed	Domain Parameter Type	Location	Motion Type	Duration
A	W	125kph	Hurricane	x1, y1, z1	Basic Motion	0:00-23:59
B	S	5kph	Rain	x2, y2, z2	Basic Motion	0:00-11:00
C	SE	10kph	Rain	x3, y3, z3	Spreading	4:00-8:53

Identifier	Direction	Speed	Domain Parameter Type	Location	Motion Type	Duration
D	S	15kph	Sleet	x4, y4, z4	Basic Motion	3:00-18:45
E	Stationary	0kph	Hail	x5, y5, z5	N/A	0:00-19:00

Table 1.

The identifier field may provide a unique identifier to each eventuality. The direction field may identify a direction of motion for the eventuality. The speed field may identify a speed of motion. The domain parameter type field may identify a type of eventuality. The location may identify a location of the eventuality in the region, and the motion type field may identify motion type (e.g., in the present example, the fact that eventuality C 306 is characterized as "spreading" due to the individual motion vectors of the component elements of the eventuality appearing to be directed away from a central location). The duration field may identify the time period in which the eventuality exists, or the time period in which the eventuality is active (e.g., moving). Some or all of these fields may be derived from the set of spatio-temporal data, though certain fields (e.g., a type field) may be derived from other data.

The eventualities may be partitioned into one or more partitions based on various factors, such as locations, type of motion, type of eventuality, direction of motion, or the like. In the instant example, the region 300 is divided into four partitions representing different geographical areas of the region. Each of the eventualities may be assigned to at least one partition. These partitions may be used to develop a document plan describing the eventualities. For example, the document plan may describe eventualities with respect to particular partitions. An example method for describing eventualities based on a partition is provided below with respect to Figure 5.

Figures 4-6 illustrate example flowcharts of the operations performed by an apparatus, such as computing system 200 of Figure 2, 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 201 of an apparatus employing an embodiment of the present invention and executed by

a processor 203 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 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 algorithms of Figures 4-6 to transform the general purpose computer into a particular 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 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. 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 depicting a method 400 for generating natural language using an input set of eventualities. The method 400 is operable to receive a set of eventualities, a region, and a particular time period. The method 400 may process these inputs to generate a document plan, and use the document plan to generate a linguistic

representation of the set of eventualities. For example, the method 400 may receive an input set of data describing one or more weather systems and motion of the weather systems, a geographical region where the weather systems are located, and a time period during which the weather systems were present, and the method 400 may
5 generate a weather forecast from this data.

At action 402, the set of eventualities, the region, and the time period are received by the method 400. In the present context, the term "region" may apply to any set of location based data associated with the set of eventualities, not just a particular geographic region. For example, in the case of weather data, the region may be a
10 geographic region of a city, county, country, or the like, but in the case of traffic data the region may be associated with a particular road or roads, in the case of network traffic the region may be associated with a particular set of nodes of the network, or in the case of medical data the region may be associated with a particular part of a patient's body. The eventualities, region, and time period may relate to a set of spatio-temporal data received
15 and processed by a motion description system. As an example, a set of eventualities may provide data about one or more weather systems and corresponding motion, the region may be a geographical region, and the particular time period may be a time of day associated with a set of spatio-temporal data from which the eventualities were derived.

At action 404, the eventualities are organized and/or partitioned. Partitioning
20 and/or organizing of the eventualities in this manner may provide for efficient description of the region of interest by grouping eventualities according to common characteristics that make for easy description. For example, weather events that occur in a particular part of the region may be grouped together for description or analysis. An example of a method for organizing and/or partitioning eventualities is described further below with
25 respect to Figure 5.

At action 406, the organized and/or partitioned eventualities are placed in a particular order. The eventualities may be ordered based on various factors, including but not limited to the type of eventuality, the partition into which the eventuality has been placed, or the importance of the eventuality as determined based on rules included in a
30 domain model 116. An example of a method for ordering the eventualities is described further below with respect to Figure 6.

At action 408, a document plan is generated using the organized, partitioned, and/or ordered eventualities. Once the eventualities are grouped and ordered, a structure may be formed for generation of one or more messages to describe the set of
35 eventualities. For example, rules for establishing messages may be applied based on a domain model 116 to convert the organized, partitioned, and/or ordered eventualities into messages. The structure of the eventualities may be used to link the messages together

into a hierarchical structure that is usable by a natural language generation system or a component thereof (e.g., a microplanner such as the microplanner 132) to generate natural language from the messages (e.g., by populating the document plan with information contained within the domain model and/or set of eventualities).

5 At action 410, a linguistic representation is generated using the document plan. As described above with respect to Figure 1, the document plan generated at action 408 may be populated with data according to the structure of the document plan and message types defined by the set of eventualities and/or domain model. The method 400 may employ a natural language generation system such as the natural language generation
10 system 108 described with respect to Figure 1 to generate said linguistic representation using the document plan.

 Figure 5 is a flow chart depicting an example of a method 500 for organization and/or partitioning a set of eventualities in accordance with example embodiments of the present invention. In order to efficiently describe the set of eventualities, embodiments of
15 the invention may organize particular eventualities based on quantity and importance. For example, if a given set of eventualities has hundreds of individual eventualities, it may not be practical to generate a set of text that describes each eventuality in detail. As such, a method such as the method 500 may be employed to organize, partition, and/or filter the eventualities to provide a concise, efficient set of eventualities for generation of
20 the linguistic representation as described with respect to Figure 4.

 At action 502, a set of eventualities represented as messages 114 is provided as input, and importance values are determined for the eventualities. The term "importance" is intended as a general term to describe analysis of the set of data to identify relevant eventualities for processing by further steps of the natural language generation process.
25 Individual eventualities or groups of eventualities may be given a particular score to rate how important description of that particular eventuality or group of eventualities is to understanding of the set of eventualities as a whole. These importance scores may be derived according to a domain model, such as the domain model 116. The importance of an eventuality may be calculated using an externally specified importance value (e.g., an
30 a priori importance value) for the eventuality type (e.g., an importance value for a particular type of weather system), multiplied by a number of spatial points in the eventuality, multiplied by the time period of the eventuality. In some embodiments, different features of the eventuality may be accorded different weights. As an example of different relative importance values, in a weather domain, certain eventualities or
35 eventuality types (e.g., a tornado or hurricane) may be deemed as more important than other eventualities (e.g., a single cloud), or a weather system that is larger, denser, or with a greater precipitation value may assigned a higher importance value than a smaller,

more dispersed, or lighter precipitation system. As another example, in the traffic domain, a car accident on a highway may be assigned a higher importance than free-flowing traffic on a back road.

At action 504, a determination is made as to whether the number of eventualities is greater than some threshold number. It may only be practical to describe a certain number of eventualities based on limitations in text output length, processor availability, or the like. As such, the method 500 may employ methods of filtering the eventualities if these are greater than some threshold number of eventualities. This threshold may be defined as a particular constant value (e.g., no more than 50 eventualities), or the number may be dynamic based on particular characteristics of the domain model and/or the set of eventualities.

At action 506, if the number of eventualities is greater than the threshold, then one or more eventualities may be filtered out. The filtering process may include dropping these eventualities entirely, or aggregating filtered eventualities and providing less detail or an aggregate description of the filtered eventualities. Eventualities may be filtered based on various criteria, including but not limited to importance, ease of aggregation, or the like. As an example, the least important eventualities may be filtered from a set of eventualities that contains greater than the threshold number of eventualities, until the set of eventualities contains less than the threshold number. As a specific example, returning to Figure 3, the eventuality B 304 might be filtered out due to its smaller size (and thus likely lower importance score) than the other eventualities present in the region, while the eventuality A 302 might be assigned a very high importance score due to the fact that hurricanes are dangerous weather systems that are highly relevant to a weather forecast.

At action 508, a determination is made as to whether the eventualities can be effectively partitioned according to a partitioning method. As described above, the partitioning may be performed in a spatial manner, such that eventualities located in similar locations are partitioned together for description, although other partitioning methods might also be employed (e.g., by type, size, importance, etc.). Spatial partitioning may be employed such that all eventualities in a particular part of the region (e.g., northern area or southern area) are grouped together. If such a partitioning is possible, the eventualities are partitioned at action 512. Otherwise, the eventualities may be treated as a single partition at action 510.

After the eventualities are partitioned at action 512, these partitioned may be ordered in a particular way at action 514. The partitions may be ordered to enhance coherence. For example, if the partitions are spatially divided into series of partitions associated with different cardinal directions (e.g., a northeastern partition, a southeastern partition, a southwestern partition, and a northwestern partition such as depicted in Figure

3), then a clockwise or counterclockwise ordering may improve coherence when describing the eventualities. The end result of the method 500 may be a partitioned set of eventualities 516, even if said partitioned set only includes a single partition.

Figure 6 is a flow chart depicting an example of a method 600 for ordering
5 eventualities as part of generating a document plan for a set of eventualities according to example embodiments of the present invention. In addition to organizing and partitioning the eventualities into groups, generation of a document plan may also include determining an appropriate order in which the eventualities (and thus eventually
10 messages associated with said eventualities) will appear. This ordering process may be employed to present eventualities in a logical, efficient order. It should be appreciated that the method 600 only describes one example embodiment for ordering of eventualities, and that additional or alternative methods and criteria could be employed for determining an order of a set of eventualities within a document plan.

At action 602, the partitioned set of eventualities 516 is received as an input and
15 time relevance is determined for each eventuality in a particular partition. Although the method 600 is not shown as looping through each partition in the figure, it should be appreciated that partitions may be processed one at a time or as a group, with the method being employed to order the eventualities within each partition, and the partitions themselves based on various factors.

20 As described above with respect to Figure 3, each eventuality may be associated with a time value or range. These time values or ranges may define the temporal period of the spatio-temporal data for which the eventuality is relevant, such as when the eventuality is present or in motion. These time values may be examined for the purposes of ordering the eventualities. For example, at action 604, a determination is made as to
25 whether a particular proportion of the eventualities (e.g., greater than 75% of all eventualities in the particular partition) apply to an entire time period, or at least a certain portion of the entire time period. If greater than the threshold proportion of eventualities within the partition apply to the entire time period, the method 600 may proceed to action 612, where these eventualities are ordered in accordance with their importance.
30 Otherwise, the method 600 proceeds to action 606.

At action 606, a determination is made as to whether the difference in importance
between the most important eventuality and the next most important eventuality is greater than a threshold value. In other words, if a particular eventuality is so much more
35 important than other eventualities in the particular partition, then it may be appropriate to indicate that the most important eventuality should be described first. If this is the case, the method 600 proceeds to action 610. Otherwise, the method 600 proceeds to action 608.

At action 608, eventualities may be ordered based on their start time, ensuring that eventualities are described in chronological order. Otherwise, if the method 600 proceeded to action 610, the most important eventuality may be described first, regardless of time, and the remaining eventualities may be ordered based on their
5 respective start time. Although action 610 describes only a single most important event being described out of chronological order, it should be appreciated that other important events (e.g., events that are greater than a particular number of standard deviations of importance than other events in the set of eventualities) could also be described out of chronological order. The end result of the method 600 is a set of ordered eventualities,
10 which may be used to generate a document plan, such as described above with respect to Figure 4.

Many modifications and other embodiments of the inventions set forth herein will 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.
15 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 descriptions and the associated drawings describe example embodiments in the context of certain example combinations of elements and/or functions, it should be appreciated
20 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 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
25 descriptive sense only and not for purposes of limitation.

THAT WHICH IS CLAIMED:

1. A method comprising:
 - receiving a set of eventualities, the set of eventualities describing at least one of a domain event and a domain state, the at least one of the domain event and the domain state derived from a set of spatio-temporal data and the set of eventualities associated with a particular region and a particular time period;
 - organizing the set of eventualities to generate a document plan; and
 - generating, using a processor, a linguistic representation of the set of eventualities using the document plan.
2. The method of claim 1, wherein organizing the set of eventualities comprises at least one of filtering out one or more eventualities, partitioning one or more of the set of eventualities into a portion of the particular region, and ordering the set of eventualities into a particular order.
3. The method of claim 1, wherein the particular region is a geographic region.
4. The method of claim 1, wherein the set of eventualities are organized according to a domain model.
5. The method of claim 4, wherein the domain model comprises a set of rules for determining an importance for one or more of the set of eventualities, and wherein organizing the set of eventualities further comprises:
 - determining the importance for one or more of the set of eventualities using the domain model; and
 - organizing the set of eventualities based on the importance.
6. The method of claim 5, wherein organizing the set of eventualities based on the importance further comprises placing a most important eventuality first in the document plan.
7. The method of claim 6, wherein the most important eventuality is placed first in the document plan in response to determining that a difference in an importance score between the most important eventuality's importance score and a next most important eventuality's importance score is greater than a threshold importance score value.

8. The method of claim 4, wherein the domain model is associated with a domain of the set of eventualities.

9. The method of claim 8, wherein the domain of the set of eventualities is at least one of weather data, traffic data, medical data, scientific data, and computer
5 network data.

10. The method of claim 2, wherein organizing the set of eventualities further comprises ordering the set of eventualities based on a start time of one or more of the set of eventualities.

11. The method of claim 1, wherein organizing the set of eventualities
10 comprises partitioning the set of eventualities into one or more portions of the region, and the set of eventualities are organized in the document plan based on a respective portion of the region into which each eventuality was partitioned.

12. The method of claim 11, further comprising ordering the partitions into a particular order to improve coherence of the document plan.

13. The method of claim 1, wherein organizing the set of eventualities
15 comprises:

attempting to separate the set of eventualities into one or more portions of the region;

determining that it is not possible to separate the set of eventualities into
20 one or more portions of the region; and

in response to determining that it is not possible to separate the set of eventualities into one or more portions of the region, organizing the document plan as a single partition.

14. The method of claim 1, wherein generating the linguistic representation
25 comprises:

generating a set of messages from the organized set of eventualities; and
conducting document planning, microplanning, and realization using the set of messages and the document plan to result in an output text.

30

15. An apparatus comprising:
a memory coupled to at least one processor; and
the at least one processor, configured to:

receive a set of eventualities, the set of eventualities describing at least one of a domain event and a domain state, the at least one of the domain event and the domain state derived from a set of spatio-temporal data and the set of eventualities associated with a particular region and a particular time period;

5 organize the set of eventualities to generate a document plan; and
 generate a linguistic representation of the set of eventualities using the document plan.

16. The apparatus of claim 15, wherein organizing the set of eventualities comprises at least one of filtering out one or more eventualities, partitioning one or more
10 of the set of eventualities into a portion of the particular region, and ordering the set of eventualities into a particular order.

17. The apparatus of claim 15, wherein the particular region is a geographic region.

18. The apparatus of claim 15, wherein the set of eventualities are organized
15 according to a domain model.

19. The apparatus of claim 18, wherein the domain model comprises a set of rules for determining an importance for one or more of the set of eventualities, and wherein the processor is further configured to organize the set of eventualities by:

 determining the importance for one or more of the set of eventualities
20 using the domain model; and
 organizing the set of eventualities based on the importance.

20. The apparatus of claim 19, wherein the processor is further configured to organize the set of eventualities based on the importance by placing a most important eventuality first in the document plan.

25 21. The apparatus of claim 19, wherein the most important eventuality is placed first in the document plan in response to determining that a difference in an importance score between the most important eventuality's importance score and a next most important eventuality's importance score is greater than a threshold importance score value.

30 22. The apparatus of claim 18, wherein the domain model is associated with a domain of the set of eventualities.

23. The apparatus of claim 22, wherein the domain of the set of eventualities is at least one of weather data, traffic data, medical data, scientific data, and computer network data.

24. The apparatus of claim 16, wherein organizing the set of eventualities
5 further comprises ordering the set of eventualities based on a start time of one or more of the set of eventualities.

25. The apparatus of claim 15, wherein organizing the set of eventualities comprises partitioning the set of eventualities into one or more portions of the region, and the set of eventualities are organized in the document plan based on a respective portion
10 of the region into which each eventuality was partitioned.

26. The apparatus of claim 25, wherein the processor is further configured to order the partitions into a particular order to improve coherence of the document plan.

27. The apparatus of claim 15, wherein the processor is further configured to organize the set of eventualities by:
15 attempting to separate the set of eventualities into one or more portions of the region;
determining that it is not possible to separate the set of eventualities into one or more portions of the region; and
in response to determining that it is not possible to separate the set of
20 eventualities into one or more portions of the region, organizing the document plan as a single partition.

28. The apparatus of claim 15, wherein the processor is further configured to generate the linguistic representation by:
25 generating a set of messages from the organized set of eventualities; and
conducting document planning, microplanning, and realization using the set of messages and the document plan to result in an output text.

29. A non-transitory computer readable storage medium comprising
30 instructions, that, when executed by a processor, configure the processor to:
receive a set of eventualities, the set of eventualities describing at least one of a domain event and a domain state, the at least one of the domain event and the domain state derived from a set of spatio-temporal data and the set of eventualities associated with a particular region and a particular time period;

organize the set of eventualities to generate a document plan; and
generate a linguistic representation of the set of eventualities using the
document plan.

5 30. The non-transitory computer readable storage medium of claim 29,
wherein instructions to organize the set of eventualities comprise at least one of filtering
out one or more eventualities, partitioning one or more of the set of eventualities into a
portion of the particular region, and ordering the set of eventualities into a particular order.

31. The non-transitory computer readable storage medium of claim 29,
wherein the particular region is a geographic region.

10 32. The non-transitory computer readable storage medium of claim 29,
wherein the set of eventualities are organized according to a domain model.

33. The non-transitory computer readable storage medium of claim 32,
wherein the domain model comprises a set of rules for determining an importance for one
or more of the set of eventualities, and wherein the non-transitory computer readable
15 storage medium further comprises instructions to configure the processor to organize the
set of eventualities by:

determining the importance for one or more of the set of eventualities
using the domain model; and

organizing the set of eventualities based on the importance.

20 34. The non-transitory computer readable storage medium of claim 33,
wherein the non-transitory computer readable storage medium further comprises
instructions to configure the processor to organize the set of eventualities based on the
importance by placing a most important eventuality first in the document plan.

35. The non-transitory computer readable storage medium of claim 33,
25 wherein the most important eventuality is placed first in the document plan in response to
determining that a difference in an importance score between the most important
eventuality's importance score and a next most important eventuality's importance score
is greater than a threshold importance score value.

36. The non-transitory computer readable storage medium of claim 32,
30 wherein the domain model is associated with a domain of the set of eventualities.

37. The non-transitory computer readable storage medium of claim 36, wherein the domain of the set of eventualities is at least one of weather data, traffic data, medical data, scientific data, and computer network data.

38. The non-transitory computer readable storage medium of claim 30,
5 wherein organizing the set of eventualities further comprises ordering the set of eventualities based on a start time of one or more of the set of eventualities.

39. The non-transitory computer readable storage medium of claim 29,
wherein organizing the set of eventualities comprises partitioning the set of eventualities
into one or more portions of the region, and the set of eventualities are organized in the
10 document plan based on a respective portion of the region into which each eventuality
was partitioned.

40. The non-transitory computer readable storage medium of claim 39, further
comprising instructions to order the partitions into a particular order to improve coherence
of the document plan.

41. The non-transitory computer readable storage medium of claim 29, further
15 comprising instructions to configure the processor to organize the set of eventualities by:
attempting to separate the set of eventualities into one or more portions of
the region;
determining that it is not possible to separate the set of eventualities into
20 one or more portions of the region; and
in response to determining that it is not possible to separate the set of
eventualities into one or more portions of the region, organizing the document plan as a
single partition.

42. The non-transitory computer readable storage medium of claim 29, further
25 comprising program instructions to generate the linguistic representation by:
generating a set of messages from the organized set of eventualities; and
conducting document planning, microplanning, and realization using the
set of messages and the document plan to result in an output text.

30

100

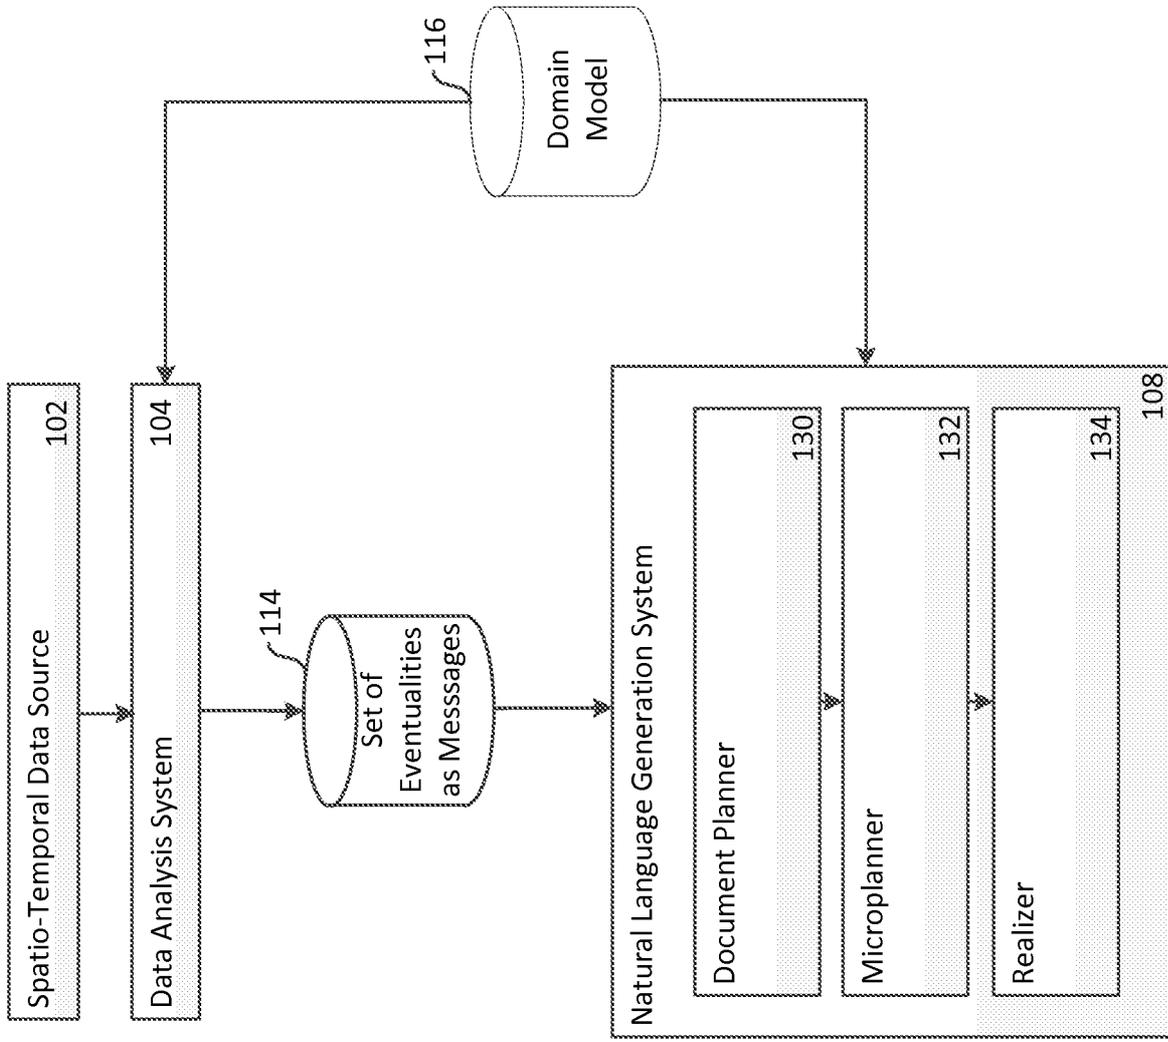


FIG. 1

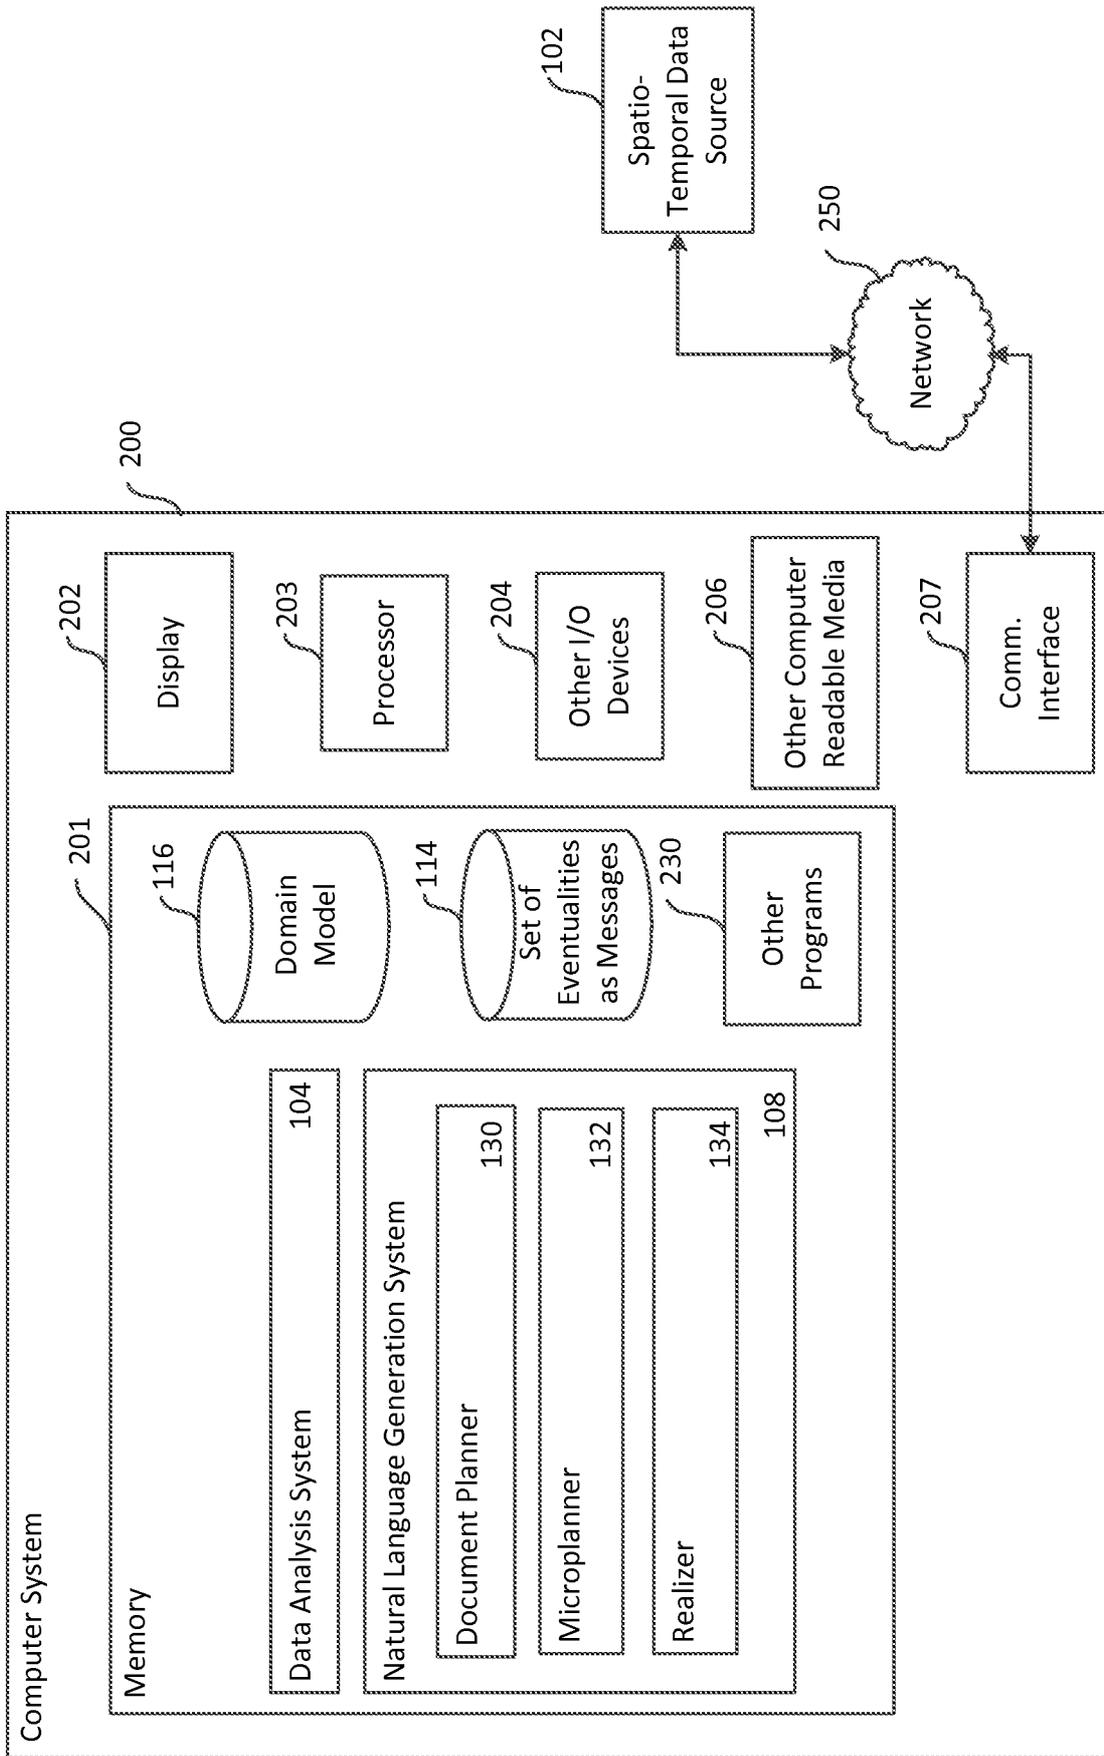


FIG. 2

300

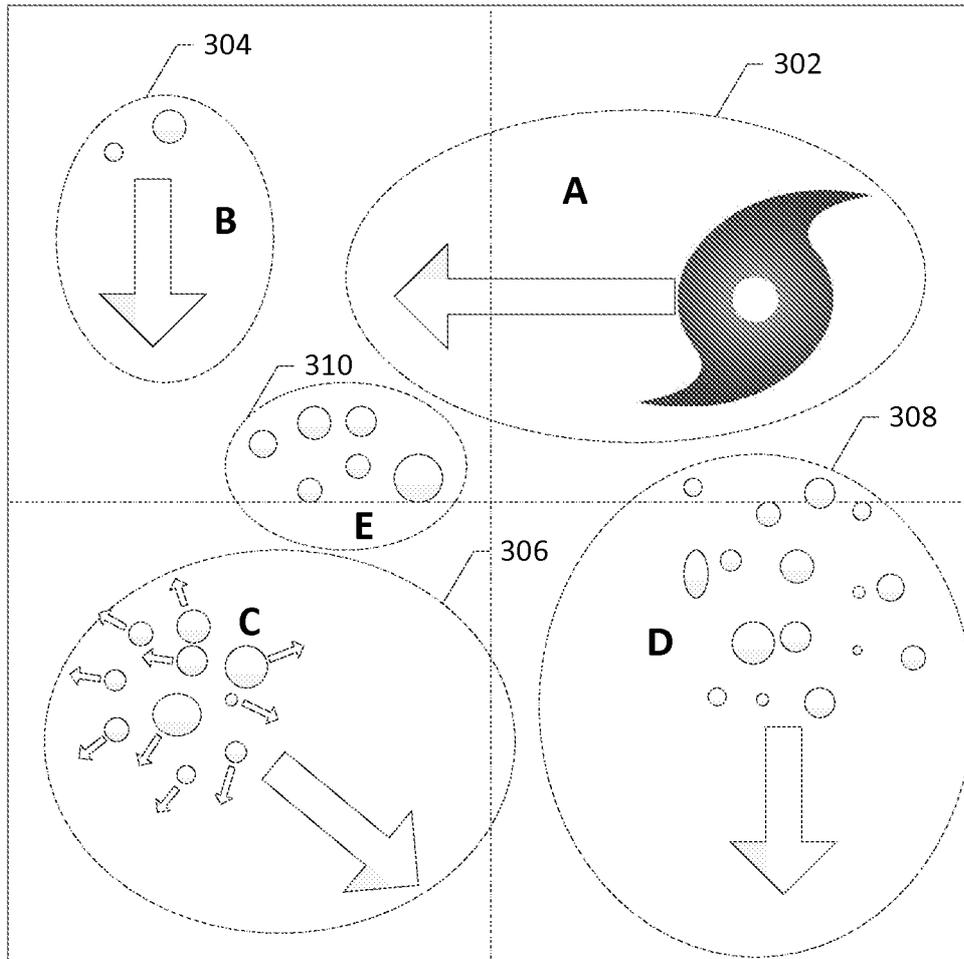


FIG. 3

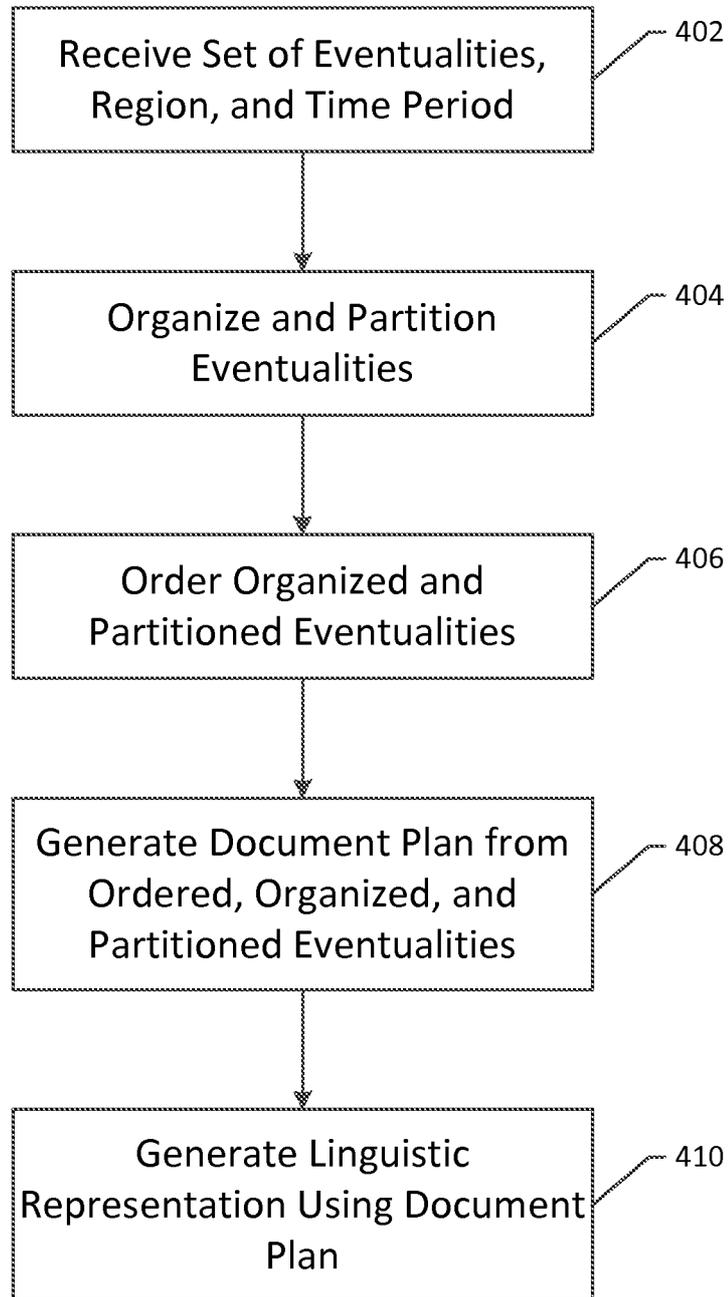


FIG. 4

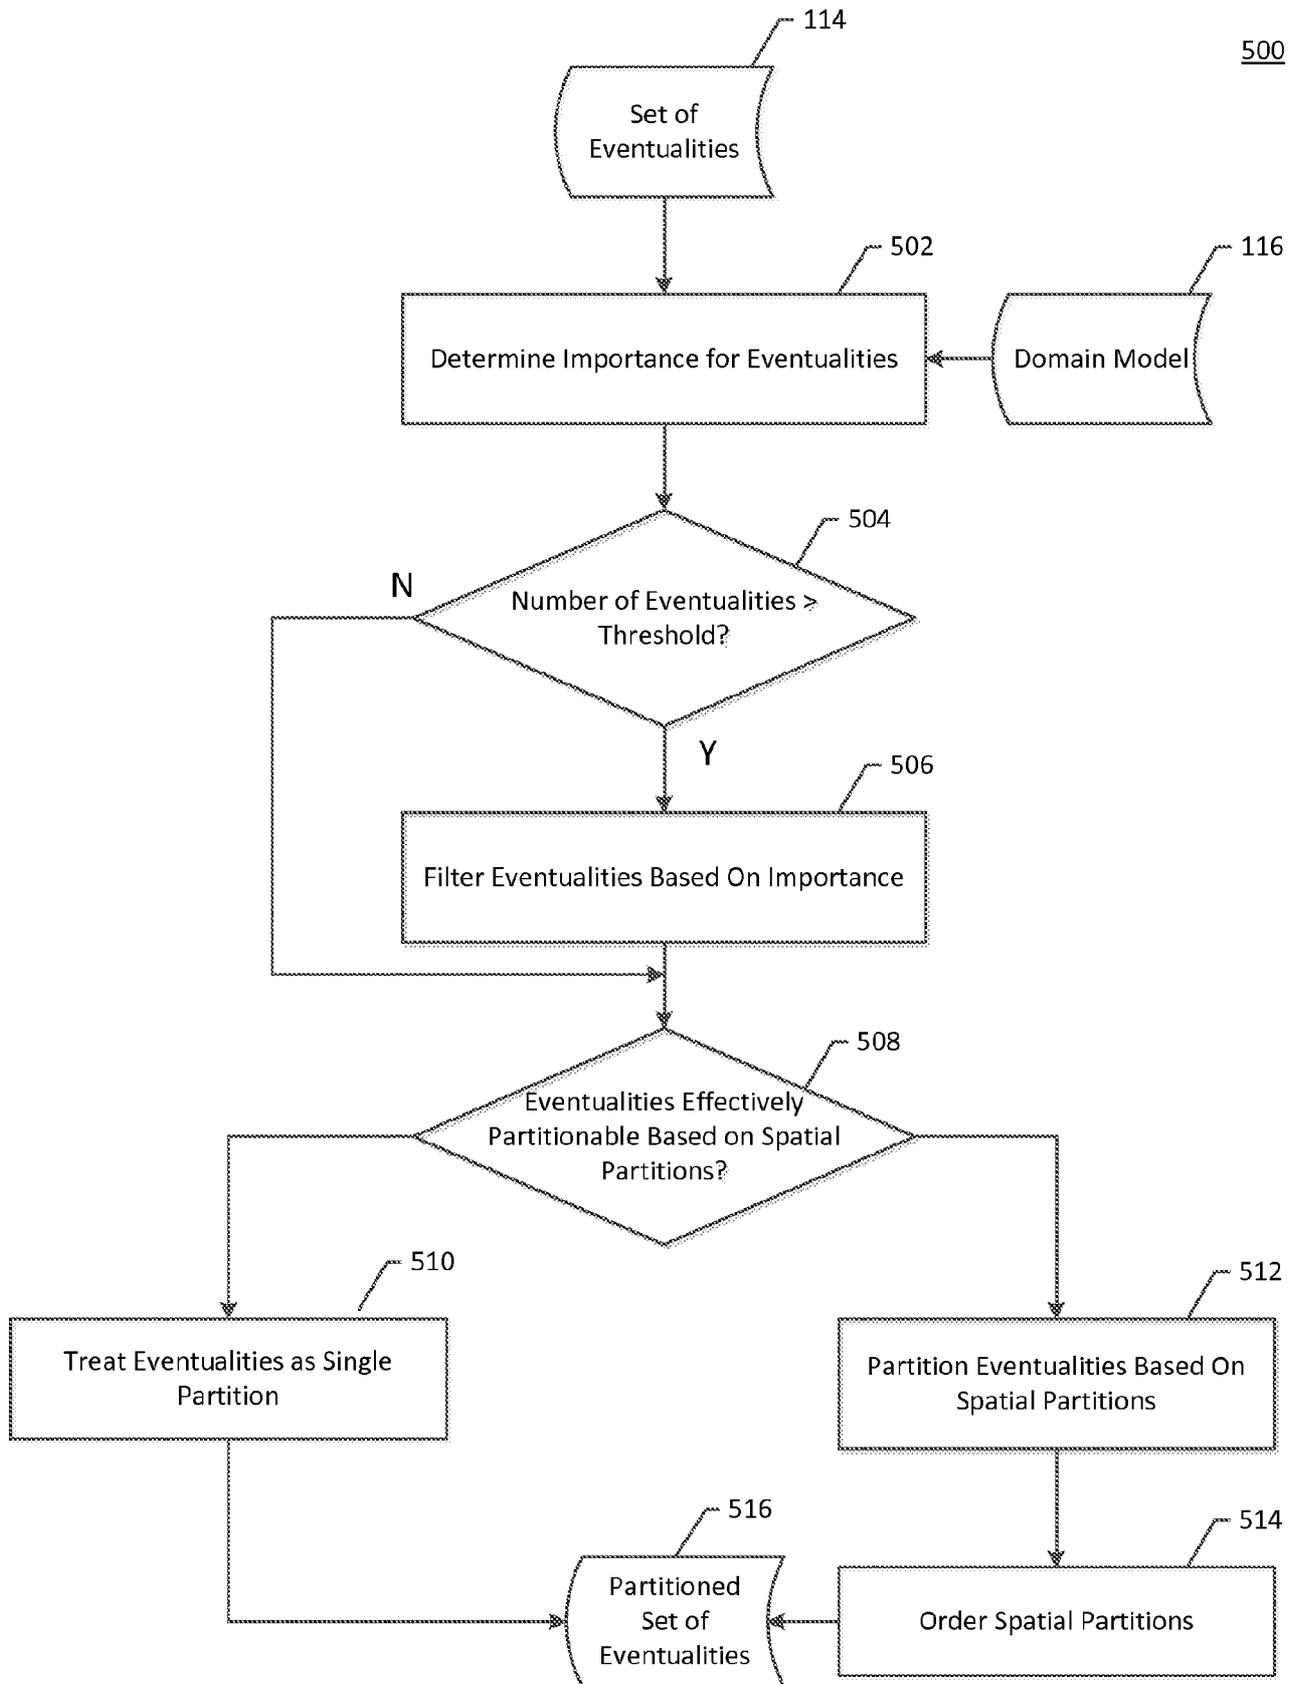


FIG. 5

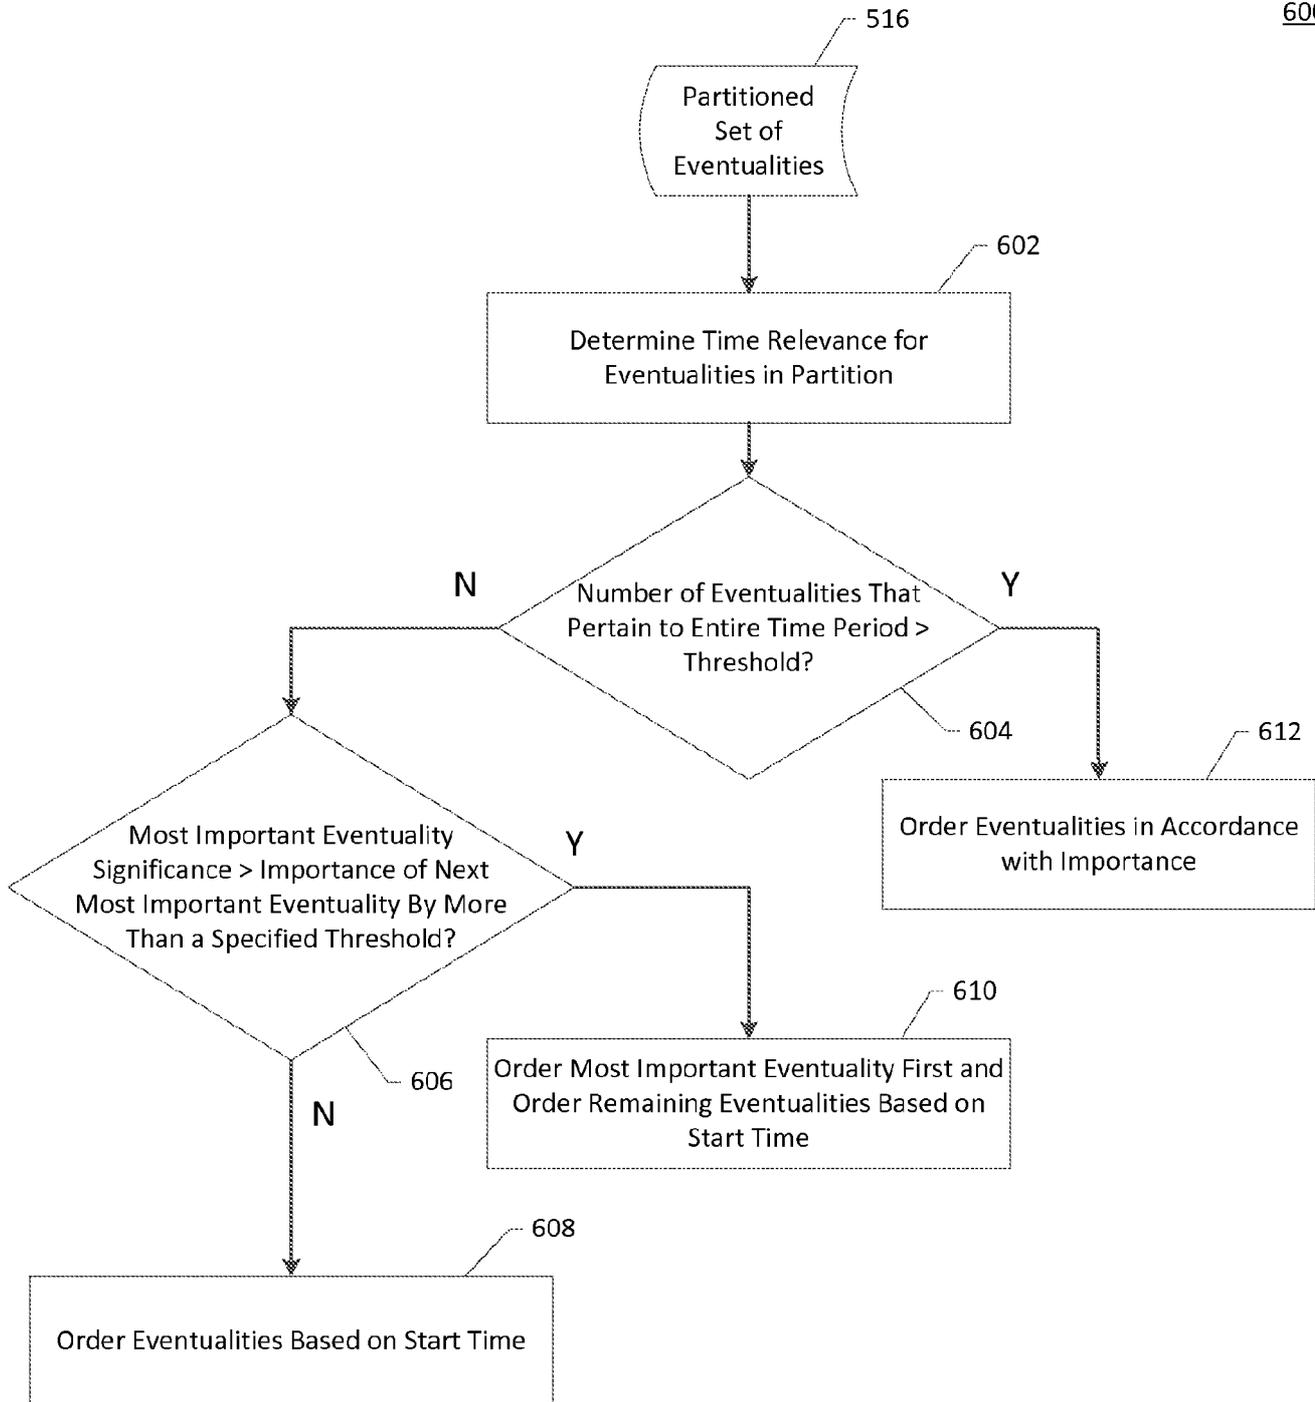


FIG. 6

INTERNATIONAL SEARCH REPORT

International application No PCT/IB2012/057774
--

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06T7/20
 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)
G06T G05B

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	ELISABETH ANDRÉ ET AL: "From visual data to multimedia presentations", GROUNDING REPRESENTATIONS: INTEGRATION OF SENSORY INFORMATION IN NATURAL LANGUAGE PROCESSING, ARTIFICIAL INTELLIGENCE AND NEURAL NETWORKS, IEEE COLLOQUIUM ON, 15 May 1995 (1995-05-15) , pages 1-3 , XP055078513 , DOI : 10.1049/ic:19950661	1-4, 8-10, 14-18, 22-24, 28-32 , 36-38,42
Y	the whole document ----- -/--	5-7 , 11-13 , 19-21 , 25-27 , 33-35 , 39-41

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

* Special categories of cited documents :

<p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier application or patent but published on or after the international filing date</p> <p>"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)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p>	<p>"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</p> <p>"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</p> <p>"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</p> <p>"&" document member of the same patent family</p>
---	---

Date of the actual completion of the international search 13 September 2013	Date of mailing of the international search report 20/09/2013
---	---

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 Huguet Serra, G
--	--

INTERNATIONAL SEARCH REPORT

International application No

PCT/IB2012/057774

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	Voelz, Dirk: "Rocco: A RoboCup Soccer Commentator System", 1999, XP55079122, Retrieved from the Internet: URL: http://www.dfki.de/~flint/papers/robocup98.pdf [retrieved on 2013-09-13]	5-7, 11-13, 19-21, 25-27, 33-35, 39-41
A	page 54 - page 56	1-4, 8-10, 14-18, 22-24, 28-32, 36-38,42
A	----- Gerd Herzog ET AL: "Combining Alternatives in the Multimedia Presentation of Decision Support Information for Real-Time Control", Designing Effective and Usable Multimedia Systems, 1998, pages 143-157, XP055079086, Retrieved from the Internet: URL: http://www.dfki.de/~flint/papers/deums98.pdf [retrieved on 2013-09-13] page 2	1-42
A	----- KOTTKE D P ET AL: "MOTION ESTIMATION VIA CLUSTER MATCHING", TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, IEEE, PISCATAWAY, USA, vol. 16, no. 11, November 1994 (1994-11), pages 1128-1132, XP000481001, ISSN: 0162-8828, DOI: 10.1109/34.334394 page 1130 - page 1131 -----	1-42