



(51) International Patent Classification:
G06F 19/00 (201 1.01)

23 10 Sawtelle Boulevard, #209, Los Angeles, CA 90066 (US).

(21) International Application Number:
PCT/US2013/057137

(74) Agent: **SHELDON, David, P.**; Christensen O'Connor Johnson Kindness PLLC, 1420 Fifth Avenue, Suite 2800, Seattle, WA 98101 (US).

(22) International Filing Date:
28 August 2013 (28.08.2013)

(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, 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, SA, 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.

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/694,171 28 August 2012 (28.08.2012) US

(71) Applicant: **THE REGENTS OF THE UNIVERSITY OF CALIFORNIA** [US/US]; 1111 Franklin Street, Oakland, CA 94607-5200 (US).

(72) Inventors: **SARRAFZADEH, Majid**; 255 Health Terrace, Anaheim, CA 92807 (US). **SUH, Myung-Kyung**; 3110 Sawtelle Boulevard, #105, Los Angeles, CA 90066 (US). **LAN, Mars**; 3245 South Sepulveda Boulevard, #203, Los Angeles, CA 90034 (US). **GHASEMZADEH, Hassan**;

(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,

[Continued on nextpage]

(54) Title: METHODS AND SYSTEMS FOR CALCULATING AND USING STATISTICAL MODELS TO PREDICT MEDICAL EVENTS

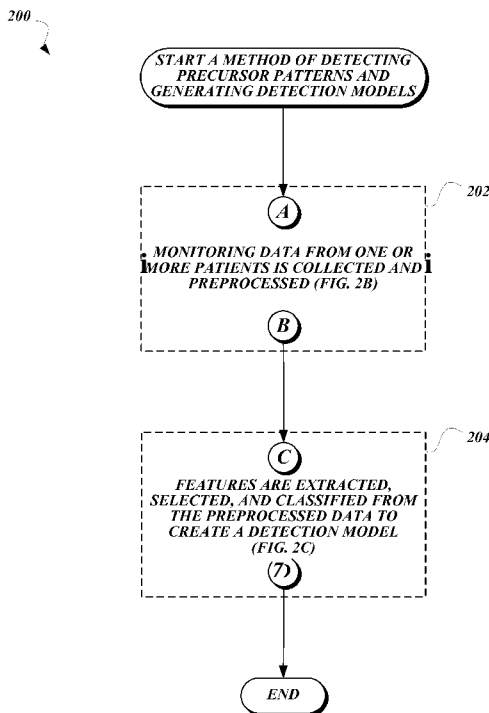


Fig.2A.

(57) Abstract: Systems and methods for generalized precursor pattern discovery that work with a wide range of biomedical signals and applications to detect a wide range of medical events are disclosed. In some embodiments, the methods and systems do not require domain-specific knowledge or significant reconfiguration based on the medical event being analyzed, hence it is also possible to discover patterns previously unknown to experts. In some embodiments, to build precursor pattern detection models, the system obtains annotated monitoring data. Positive and negative segments are extracted from the annotated monitoring data, and are preprocessed. Features are extracted from the preprocessed segments, and selected features are chosen from the extracted features. The selected features are classified to create the precursor pattern detection model. The precursor pattern detection model may then be used in real time to detect occurrences of the medical event of interest.

WO 2014/036173 A1

UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TI,
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, KM, ML, MR, NE, SN, TD, TG).

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

METHODS AND SYSTEMS FOR CALCULATING AND USING STATISTICAL
MODELS TO PREDICT MEDICAL EVENTS

CROSS-REFERENCE(S) TO RELATED APPLICATION(S)

5 This application claims the benefit of U.S. Provisional Application No. 61/694171, filed August 28, 2012, the entire disclosure of which is hereby incorporated herein by reference for all purposes.

STATEMENT OF GOVERNMENT LICENSE RIGHTS

10 The inventions disclosed herein were made with Government support under Grant No. T15 LM07356, awarded by NIH/National Library of Medicine Medical Informatics Training Program. The Government has certain rights in the disclosed inventions.

BACKGROUND

15 Recent advancements in sensor and wireless communication technologies have opened up many opportunities to acquire biomedical signals at a very low cost. Many sensors are compact enough to be worn by the subjects, and can continuously gather data for a prolonged period. These technologies have quickly found their natural applications in health care in the form of eHealth and telemedical systems.

20 Most of the early eHealth systems focus on remote monitoring and abnormality detection. While providing a low-cost and convenient way for health care personnel to monitor the well-being of patients, the majority of these systems are essentially infrastructures for data collection and storage. One of the main goals of the next generation of eHealth systems is to mine and analyze the sensor data for so-called
25 "precursor patterns." These patterns are highly correlated to an ensuing medical condition or clinical episode. Using precursor patterns to predict medical events will be important for the next generation of eHealth systems.

Thus far, there have been several studies dedicated to identifying precursor patterns with a varying degree of success. However, one common problem with these
30 studies is the requirement of domain-specific knowledge to develop their discovery algorithms. Also, most of the algorithms require prior knowledge of the duration of the precursor patterns, as well as the time the patterns are likely to occur relative to the clinical episode. Consequently, it is often impossible to apply a given precursor pattern

discovery algorithm to a different medical condition without significant modification or degradation in performance.

SUMMARY

5 This summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This summary is not intended to identify key features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter.

10 In some embodiments, a computer-implemented method of generating a precursor pattern detection model for detecting an occurrence of a given type of medical event is provided. A computing device obtains annotated monitoring data, wherein the annotated monitoring data represents measurements of physical conditions of at least one patient. A computing device extracts a set of positive segments and a set of negative segments from the annotated monitoring data, wherein each positive segment is associated with an
15 occurrence of a medical event of the given type, and wherein each negative segment is not associated with an occurrence of the medical event of the given type. A computing device executes a time filter on the set of positive segments to generate a set of filtered positive segments. A computing device executes a time filter on the set of negative segments to generate a set of filtered negative segments. A computing device extracts a
20 set of extracted positive features from the set of filtered positive segments and a set of extracted negative features from the set of filtered negative segments. A computing device selects a set of selected positive features from the set of extracted positive features and a set of selected negative features from the set of extracted negative features.

25 A computing device calculates a precursor pattern detection model based on the set of selected positive features and the set of selected negative features, and a computing device stores, on a nontransitory computer-readable medium, the precursor pattern detection model. In some embodiments, a nontransitory computer-readable medium having computer-executable instructions stored thereon is provided wherein the instructions, in response to execution by one or more processors of one or more
30 computing devices, cause the one or more computing devices to perform such a method. In some embodiments, a system comprising one or more computing devices configured to perform such a method is provided.

In some embodiments, a method of detecting an occurrence of a given type of medical event is provided. A precursor pattern detection model is retrieved from a precursor pattern data store. One or more monitoring devices are used to collect data from a patient. The collected data is analyzed using the precursor pattern detection
5 model. In response to the analysis indicating that a precursor pattern has been detected in the collected data, an alert is caused to be presented to a caregiver. In some embodiments, a computing device configured to perform such a method is provided.

In some embodiments, a nontransitory computer-readable medium having computer-executable instructions stored thereon is provided, wherein the instructions, in
10 response to execution by one or more processors of a computing device, cause the computing device to perform such a method.

DESCRIPTION OF THE DRAWINGS

The foregoing aspects and many of the attendant advantages of embodiments of
15 the subject matter disclosed herein will become more readily appreciated as the same become better understood by reference to the following detailed description, when taken in conjunction with the accompanying drawings, wherein:

FIGURE 1 is a block diagram that illustrates an exemplary embodiment of a precursor pattern analysis system according to various aspects of the present disclosure;

20 FIGURES 2A-2C are a flowchart that illustrates an exemplary embodiment of a method of detecting precursor patterns and generating detection models according to various aspects of the present disclosure;

FIGURES 3A-3B are a flowchart that illustrates an exemplary embodiment of a top-down and bottom-up time filter procedure according to various aspects of the present
25 disclosure;

FIGURE 4 is a flowchart that illustrates an exemplary embodiment of a method of using a precursor pattern detection model to detect medical events according to various aspects of the present disclosure;

FIGURE 5 is a table that illustrates various examples of classifiers used in various
30 embodiments of the present disclosure;

FIGURE 6 is a chart that illustrates prediction accuracy obtained for three datasets tested using an embodiment of the present disclosure;

FIGURE 7 is a chart that illustrates false-positive rates of the three datasets tested using the embodiment of the present disclosure;

FIGURES 8A and 8B are charts that illustrate systolic and diastolic blood pressure and weight readings from a patient over two separate 7-day periods;

5 FIGURE 9 is a chart that illustrates accuracy, sensitivity, and specificity for various classifiers and a daily weight change value in a test of an embodiment of the present disclosure;

FIGURE 10 is a chart that illustrates Receiver Operating Characteristic (ROC) curves for the various classifiers and the daily weight change value in the test of the
10 embodiment of the present disclosure; and

FIGURE 11 is a block diagram that illustrates aspects of an exemplary computing device appropriate for use with embodiments of the present disclosure.

DETAILED DESCRIPTION

15 The present disclosure describes systems and methods for generalized precursor pattern discovery that work with a wide range of biomedical signals and applications. In some embodiments, the methods and systems do not require domain-specific knowledge or significant reconfiguration based on the medical event being analyzed, hence it is also possible to discover patterns previously unknown to experts.

20

SYSTEM OVERVIEW

FIGURE 1 is a block diagram that illustrates an exemplary embodiment of a precursor pattern analysis system 100 according to various aspects of the present disclosure. As illustrated, the precursor pattern analysis system 100 includes three tiers: a
25 patient monitoring system 102, a patient data processing system 104, and a precursor pattern detection system 106. Each of the systems 102, 104, 106 may include one or more computing devices configured to provide the engines, data stores, and/or services illustrated and described herein as part of the respective system 102, 104, 106.

In some embodiments, the patient monitoring system 102 is configured to collect
30 particular measurements from the patient 90, and to transmit the collected information to the patient data processing system 104 for processing and storage. The precursor pattern detection system 106 is configured to access the collected information stored by the patient data processing system 104, to process the collected information to detect one or

more precursor patterns, and to provide precursor pattern detection models to the patient data processing system 104 for storage. In some embodiments, the patient monitoring system 102 may be further configured to use the precursor pattern detection models to monitor patients for occurrences of the precursor patterns.

5 In some embodiments, the patient monitoring system 102 may be located at the patient's home (or otherwise placed in a location easily accessible to the patient), while the patient data processing system 104 and the precursor pattern detection system 106 are placed in a central location such that multiple patient monitoring systems 102 in multiple **locations** may share the resources of a single patient data processing system 104 and/or
10 precursor pattern detection system 106. In such embodiments, the components of the precursor pattern analysis system 100 may communicate with each other using any suitable communication **technology**, such as wired or wireless Internet **communication**, direct dial telephone or ISDN communication, **and/or** any other suitable communication technology. In some embodiments, the patient monitoring system 102, the patient data
15 processing system 104, **and/or** the precursor pattern detection system 106 may be provided by a single computing device, or by multiple computing devices each located on the same network, thus obviating the need for long-distance communication between the elements of the precursor pattern analysis system 100. Though FIGURE 1 illustrates various components as being included in either the patient monitoring system 102, the
20 patient data processing system 104, or the precursor pattern detection system 106, in some embodiments, the illustrated components may be moved from one of the illustrated systems 102, 104, 106 to another of the illustrated systems 102, 104, 106 or split between two or more of the illustrated systems 102, 104, 106; two or more of the illustrated systems 102, 104, 106 may be combined into a single system; and/or one of the illustrated
25 systems 102, 104, 106 may be broken into multiple separate systems.

As illustrated, the patient monitoring system 102 includes one or more physical data sensor devices 108 and one or more interactive sensor devices 118 for **collecting** measurements from a patient 90. The physical data sensor devices 108 are physical devices that may automatically perform a measurement of a physical value, and may
30 cause the detected physical value to be automatically transmitted to the patient data processing system 104 without further interaction from the patient 90. For example, the physical data sensor devices 108 may include, but are not limited to, an electronic sphygmomanometer, an electronic thermometer, an **electronic** heart rate monitor, a

pedometer, an electronic scale, an electroencephalograph machine, an electrocardiogram machine, a motion sensor, and/or the like. The interactive sensor devices 110 are physical devices that collect information entered by the patient 90 (or by a caregiver operating the patient monitoring system 102 on behalf of the patient 90). In some
5 embodiments, an interactive sensor device 110 may be configured to present a prompt to the patient 90 to instruct the patient 90 to perform a measurement with a non-connected medical device or answer a survey question, and may subsequently allow the patient 90 to enter the result of the measurement or the answer the survey question. The interactive sensor device 110 then sends the collected information to the patient data processing
10 system 104. In some embodiments, the interactive sensor devices 110 may include a desktop computing device executing a stand-alone application or presenting a web site, a mobile computing device such as a smart phone, a tablet computing device, a custom computing device configured to present an interface to the patient monitoring system 102, and/or the like.

15 As illustrated, the patient data processing system 104 includes a data collection engine 112. In general, the word "engine" (used interchangeably with the word application or module), as used herein, refers to logic embodied in hardware or software instructions, which may be written in a programming language, such as C, C++, COBOL, JAVA™, PHP, Perl, HTML, CSS, JavaScript, VBScript, ASPX, Microsoft .NET™,
20 and/or the like. An engine may be compiled into executable programs or written in interpreted programming languages. Software engines or applications may be callable from other engines or from themselves. Generally, the engines or applications described herein refer to logical modules that can be merged with other engines or applications, or can be divided into sub-engines. The engines or applications may be stored in any type
25 of computer-readable medium or computer storage device and be stored on and executed by one or more general purpose computers, thus creating a special purpose computer configured to provide the engine or application.

In some embodiments, the data collection engine 112 may provide a web service or other appropriate application programming interface (API) with which the patient
30 monitoring system 102 communicates, such that collected data may be received by the data collection engine 112 via the API. In some embodiments, the data collection engine 112 may generate a web-based interface to be presented to the patient 90 by the patient monitoring system 102 and into which the collected data may be directly entered.

In some embodiments, the data collection engine 112 may cause the raw collected data to be stored, while in some embodiments, the data collection engine 112 may perform filtering or other data cleaning tasks before causing the collected data to be stored. Further functionality of the data collection engine 112 is described below.

5 The illustrated patient data processing system 104 also includes a patient data store 114, and a precursor pattern data store 116. As understood by one of ordinary skill in the art, a "data store" as described herein may be any suitable device configured to store data for access by a computing device. One example of a data store is a highly reliable, high-speed relational database management system (DBMS) executing on one or
10 more computing devices and accessible over a high-speed packet switched network. However, any other suitable storage technique and/or device capable of quickly and reliably providing the stored data in response to queries may be used, and the computing device may be accessible locally instead of over a network, or may be accessible over some other type of suitable network or provided as a cloud-based sendee. A data store
15 may also include data stored in an organized manner on a storage medium 1108, as described further below. One of ordinary skill in the art will recognize that separate data stores described herein may be combined into a single data store, and/or a single data store described herein may be separated into multiple data stores, without departing from the scope of the present disclosure.

20 In some embodiments, the patient data store 114 is configured to store the data collected from the patient 90 by the patient monitoring system 102. In some embodiments, the patient data store 114 may also be configured to store other information related to patients, such as demographic information, associated health care provider information, login and password information, and/or the like. In some
25 embodiments, the patient data store 114 may be configured to store data collected from the patient 90 as well as data that has been processed by the precursor pattern detection system 106. In some embodiments, the precursor pattern data store 116 is configured to store precursor pattern detection models generated by the precursor pattern detection system 106 based on the collected patient data. The precursor pattern detection models
30 stored in the precursor pattern data store 116 may be retrieved by the patient monitoring system 102 in order to monitor for occurrences of the precursor patterns in new data collected from the patient 90. In some embodiments, the precursor pattern data store 116 may also store other information associated with the precursor pattern detection models,

such as **information** about associated medical conditions, instructions to be executed upon detection of a given precursor pattern, **and/or** the like. Further aspects of the patient data store 114 and the precursor pattern data store 116 are discussed below.

As illustrated, the precursor pattern detection system 106 includes a signal preprocessing engine 120, a feature extraction engine 122, a feature selection engine 124, and a classification engine 126. In some embodiments, the signal preprocessing engine 120 is configured to preprocess collected data to place it in a format suitable for further processing by the other components of the precursor pattern detection system 106. The preprocessing may include identifying and/or removing out-of-range, missing, or otherwise invalid data; normalizing collected data received from different patient monitoring systems 102 or sensor devices to have matching scales and/or frequencies; **and/or** the like. In some embodiments, the preprocessing engine is configured to extract positive segments and negative segments from the collected data, wherein the positive segments are associated with medical events to be predicted, and the negative segments are not associated with medical events to be predicted. In some embodiments, the preprocessing engine is also configured to apply time filters to the extracted segments, such that the segments are filtered down to their most relevant portions with respect to significance in detecting the precursor pattern. Further operations of the signal preprocessing engine 120 are discussed below.

In some embodiments, the feature extraction engine 122 is configured to extract features from the segments extracted by the signal preprocessing engine 120. As discussed in further detail below, the features are a simplified representation of some aspect of the data of the segment, such as a mathematical property (standard deviation, mean, **and/or** the like), a transformation (a wavelet transform and/or the like), and/or any other simplified representation of the data. In some embodiments, the feature selection engine 124 is configured to select a set of relevant features extracted by the feature extraction engine 122. In some embodiments, the classification engine 126 is configured to execute a classification algorithm over the features selected by the feature selection engine 124, and to thereby generate the precursor pattern detection model. Further details of **the** functionality of each of these engines are **provided below**.

DETECTING PRECURSOR PATTERNS

FIGURES 2A-2C are a flowchart that illustrates an exemplary embodiment of a method of detecting precursor patterns and generating detection models according to various aspects of the present disclosure. From a start block, the method 200 proceeds to a set of method steps 202 defined between a start terminal ("terminal A") and an exit terminal ("terminal B"), wherein monitoring data from one or more patients is collected and preprocessed.

From terminal A (FIGURE 2B), the method 200 proceeds to block 206, where a patient monitoring system 102 collects data from one or more patients 90 using one or more sensor devices. The one or more sensor devices may include one or more physical data sensor devices 108, such as sphygmomanometers, electroencephalograph machines, electrocardiogram machines, thermometers, scales, and/or the like; one or more interactive sensor devices 110, such as a computing device configured to collect survey responses from the patient 90, a mobile device configured to collect information input by the patient 90, and/or the like; and/or any other suitable sensor device. Next, at block 208, the patient monitoring system 102 transmits the collected data to a data collection engine 112 of a patient data processing system 104. As described above, the collected data may be transmitted to the patient data processing system 104 via the internet, a modem, a leased line, or via any other suitable communication technology. As also described above, in some embodiments the data collection engine 112 may provide a web service or other suitable API to receive the data from the patient monitoring system 102.

At block 210, the collected data is annotated to indicate the occurrence of one or more medical events. In some embodiments, this annotation may occur automatically by the data collection engine 112 or the sensor devices themselves. For example, in some embodiments, an annotation may be made automatically in the collected data by an ICU monitor upon detecting that a vital sign being tracked by the ICU monitor has crossed a predetermined threshold. In some embodiments, this annotation may be added manually by an expert reviewing the collected data. For example, if a patient 90 is being monitored for sleep apnea and the collected data includes a polysomnogram, a reviewing physician may read the polysomnogram and manually annotate the record to indicate the sleep apnea events. At block 212, the data collection engine 112 stores the annotated collected data in a patient data store 114.

In some embodiments, the actions illustrated in blocks 206-212 and described above may be performed multiple times to gather a set of data to be analyzed for task optimization. The repetition of the actions in blocks 206-212 may be performed with a single patient over a period of time, or may be performed using data from many patients to provide greater statistical rigor for the conclusions drawn. In some embodiments, the data to be analyzed by the remainder of the method 200 may be collected using techniques other than those described above.

At block 214, a signal preprocessing engine 120 of a precursor pattern detection system 106 extracts a set of positive segments from the annotated collected data stored in the patient data store 114, and at block 216, the signal preprocessing engine 120 extracts a set of negative segments from the annotated collected data. The positive segments are segments of the annotated collected data associated with instances of the medical event during which precursor patterns are believed to occur. The negative segments are segments of the annotated collected data which are not associated with instances of the medical event and so are not expected to include the precursor patterns. In some embodiments, the positive segments may be extracted by taking a segment of a predetermined length before an annotated medical event. The negative segments may likewise be extracted by taking a segment of the same predetermined length from a portion of the annotated collected data that is distant from any annotated medical event, or from a portion of the annotated collected data associated with a patient that did not exhibit the medical event.

The method 200 then proceeds to procedure block 218, where the signal preprocessing engine 120 applies top-down and bottom-up time filters to the segments to produce a set of filtered positive segments and a set of filtered negative segments. The time filters help remove portions of the segments that do not significantly contribute to the distinctiveness of the segments, thereby increasing the strength of the detection functionality while also reducing the amount of data to be processed. An example top-down and bottom-up time filter procedure is illustrated in FIGURES 3A-3B, and is discussed further below. The method 200 then proceeds to terminal B.

From terminal B (FIGURE 2A), the method 200 proceeds to a set of method steps 204 defined between a start terminal ("terminal C") and an exit terminal ("terminal D"), wherein features are extracted, selected, and classified from the preprocessed data to create a detection model. From terminal C (FIGURE 2C), the

method 200 proceeds to block 220, where a feature extraction engine 122 of the precursor pattern detection system 106 determines a set of extracted positive features for the filtered positive segments. Each feature is a high-level, succinct way of representing or summarizing the original data in a simplified form. For example, in some embodiments, a feature may be extracted by computing a standard deviation of data points within the segment, thus representing an aspect of multiple data points with a single value. As another example, a feature for a segment may be extracted by determining a delta of values within the segment compared to values within a previous segment.

As yet another example, features could be extracted that encompass both spectral and temporal information of the signal. The spectral information is generic and widely applicable to many types of signals and many types of medical events, while the temporal information can help identify the time and length of the precursor pattern. In some embodiments, such features that encompass both spectral and temporal information of the signal may be extracted by performing a wavelet transform, where the signal is represented using orthonormal function basis called a mother wavelet. A Haar wavelet transform (or Daubechies-2) may be particularly suitable for extracting features, as it has a low $O(n)$ complexity and has been shown to work well with time series data. The wavelet coefficients may be computed at multiple resolutions, each at half of the scale of the previous one. These coefficients are then used as the extracted features.

In some embodiments, multiple features may be extracted for a single segment, such as a mean and a standard deviation, a standard deviation and a delta, a standard deviation and a wavelet transform, or any other combination of two or more features. Likewise, multiple features may be extracted for a single segment based on multiple signals included in the segment. For example, feature extraction may extract a standard deviation of an electrocardiogram signal, and may extract a delta for a body weight signal, both for the same segment.

Next, at block 222, a feature selection engine 124 of the precursor pattern detection system 106 determines a strength for each feature of the set of extracted positive features, and at block 223, the feature selection engine 124 selects a set of features from the set of extracted positive features based on the determined strengths of the features using a wrapper algorithm. Performing feature selection helps reduce the amount of data to be processed to make classification more efficient and accurate. For example, a one minute segment on a 128Hz signal generates a total of 7680 coefficients

when performing a Haar wavelet transform. Building a classification model using every available coefficient would not only be slow, but would also be likely to result in overfitting. Only a small portion of the coefficients are likely correlate to the medical event and thus constitute the precursor pattern.

5 Generally, there are three types of feature selection algorithms: Wrapper, Embedded, and Hybrid. Wrapper algorithms may be the most appropriate, as they select features based on aspects of the features, such as correlation, relevance, and redundancy, and are therefore model-agnostic. In some embodiments, Information Gain may be used to calculate the feature strength and form the basis for feature selection. The Information
 10 Gain (*IG*) for a feature J_i given a set of training samples S is:

$$IG(S_x, f_i) = H(S_x) - H(S_x | f_i)$$

$H(S)$ is the information entropy of S :

$$H(S) = - \sum_{i=1}^n p(S_i) \log p(S_i)$$

The conditional entropy $H(S_x | f_i)$ is therefore:

15
$$H(S_x | f_i) = \sum_{x \in S_x} p(x, f_i) \log \frac{p(f_i)}{p(x, f_i)}$$

In other words, $IG(S_x, f_i)$ is the change in entropy if f_i is known in advance. A feature with small *IG* is considered less relevant and can be discarded without weakening the classification model. In some embodiments, features having a strength below a predetermined threshold may be discarded. In some embodiments, a top percentage or
 20 absolute number (e.g., top-100) of features as ranked by strength may be retained, and the rest of the features may be discarded.

At block 224, the feature extraction engine 122 determines a set of extracted negative features for the filtered negative segments; at block 225, the feature selection engine 124 determines a strength for each feature of the set of extracted negative features;

and at block 227, the feature selection engine 124 selects a set of negative features from the set of extracted negative features based on the determined strengths of the features using a wrapper algorithm. In blocks 224, 225, and 227, the method 200 performs actions similar to those discussed above with respect to blocks 220, 222, and 223, respectively, but over the filtered negative segments as opposed to the filtered positive segments. Because these actions are similar to those discussed above with respect to blocks 220, 222, and 223, the actions performed at blocks 224, 225, and 227 are not described further for the sake of brevity.

The method 200 then proceeds to block 228, where a classification engine 126 of the precursor pattern detection system 106 constructs a precursor pattern detection model based on the selected set of features. While constructing the model, it is important to separate the training data from the testing data to prevent model overfitting and overly optimistic results. However, if the number of positive segments is limited due to the rare nature of the medical event within the collected data, it is also possible to perform k-fold validation using the same number of positive and negative segments. The method 200 is classifier agnostic, and so any type of classifier known to one of skill in the art may be used to construct the precursor pattern detection model from the selected set of features. Some examples of suitable classifier algorithms include, but are not limited to, Naive Bayes, Bayes Network, Logistic Regression, C4.5 Decision Tree, Support Vector Machine (SVM), Voting Feature Interval (VFI), Artificial Neural Network (ANN), Nearest Neighbor (kNN), and Ripple-Down Rule Learner (RIDOR). Though the use of these classifier algorithms are known to one of skill in the art, brief descriptions of some of these classifier algorithms have nevertheless been included in FIGURE 5.

After the precursor pattern detection model is constructed, the method 200 proceeds to block 230, where the classification engine 126 stores the precursor pattern detection model in a precursor pattern data store 116 of the patient data processing system 104. Information stored in the precursor pattern data store 116 along with the precursor pattern detection model may include an identification of the medical event to be detected by the precursor pattern detection model, an identification of the signals considered by the precursor pattern detection model, a set of instructions to be executed upon detecting the precursor pattern, and/or the like. The method 200 then proceeds to terminal D, and from terminal D (FIGURE 2A) to an end block, where the method 200 terminates.

FIGURES 3A-3B are a flowchart that illustrates an exemplary embodiment of a top-down and bottom-up time filter procedure according to various aspects of the present disclosure. The illustrated procedure 300 is an example of a procedure suitable for execution at block 218 in the method 200 described above. In some embodiments, all segments of the annotated collected data are lined up and considered together. For example, if segments are five days long, all of the first days of the segment are considered together, all of the second days of the segment are considered together, and so on, such that a determination can be made regarding the strength of portions of a generic segment as opposed to an individual actual segment.

The procedure 300 begins at block 302, where a feature strength of a segment is determined. Feature strength may be determined using any suitable technique for determining a correlation between the feature and the classification of the segment. For example, in some embodiments, an Information Gain (IG) value is computed and used as the feature strength. The procedure 300 then proceeds to block 304, where the segment is subdivided into two or more subsegments. In some embodiments, the segment may be subdivided based on data granularity. For example, if there is one data point per day, a multi-day segment may be subdivided into days. In some embodiments, the subdivisions may include multiple data points if data granularity is high compared to segment length. For example, if data is sampled at 60Hz and the segment is a minute long, the segment may be subdivided into 60-data-point subdivisions (one subdivision per second).

At block 306, each subsegment is removed from the segment in turn, and a feature strength is determined for the remaining portion of the segment upon removing each subsegment. Next, at block 308, a subsegment is determined that, upon removal, had the largest positive impact on the feature strength of the remaining portion of the segment, and at block 310, the determined subsegment is removed from the segment. For example, if the segment is five days long and has been subdivided into five single-day subdivisions, removing day 1 may have no impact on the feature strength of the subsegment consisting of days 2, 3, 4, and 5 compared to the feature strength of the segment as a whole. Meanwhile, removing day 3 may increase the feature strength of the subsegment consisting of days 1, 2, 4, and 5 compared to the feature strength of the segment as a whole. Accordingly, in this example day 3 would be removed, and the remaining segment would consist of days 1, 2, 4, and 5. The intuition is that portions of the segment that do not significantly contribute to the feature strength can be removed to

both streamline processing and to increase the predictive strength of the remaining portions of the segment.

In some embodiments, division into **subsegments** may include separating individual **signals or features**. For example, in the five-day example **segment** described above, there may be two features computed for the segment. The two features may be extracted from the same signal (e.g., a standard deviation and a mean of a time series), or could be extracted from different signals (e.g., a standard deviation of a time series of a first EEG signal, and a standard deviation of a time series of a second EEG signal). In such embodiments, each extracted feature may be divided into a separate **subsegment**, such that **when** the five-day example segment is broken into day-long subsegments, ten subsegments may be created (one for each of the two features for each day).

Once the determined subsegment is removed from the segment, the procedure 300 proceeds to a decision block 312, where a **determination** is made as to whether top-down filtering should continue on the segment. In some embodiments, top-down filtering is continued if the strength of the segment is not yet above a predetermined threshold strength. In some embodiments, top-down filtering is continued even if the strength of the segment is above a predetermined threshold strength, as long as the strength of the segment is still increasing as a result of the execution of the filter. If the result of the determination at decision block 312 is YES, the procedure 300 returns to block 302, and another iteration of the top-down filtering steps is performed on the segment to remove another subsegment. Otherwise, if the result of the determination at decision block 312 is NO, the procedure 300 advances to a continuation terminal ("terminal X").

From terminal X (FIGURE 3B), the procedure 300 proceeds to block 314, **where** the segment is divided into two or more subsegments of a given size. The size of the subsegments may be determined similarly to the subsegments discussed above, and the subsegments at block 314 may be the same size as the subsegments in the top-down filtering discussed above, or may be a different size. At block 316, a feature strength is determined for each subsegment. Next, at block 318, each subsegment is combined with another subsegment that has a least negative impact on the feature strength. This bottom-up combination of subsegments further reduces the number of features present without negatively impacting the overall predictive strength of the segment.

The procedure 300 proceeds to a decision block 320, where a determination is made as to whether the current subsegment is the last remaining subsegment. If the result

of the determination at decision block 320 is YES, the procedure 300 advances to a continuation terminal ("terminal Z"). Otherwise, if the result of the determination at decision block 320 is NO, the procedure 300 advances to block 322, where, if a strength of the subsegment is below a threshold strength, the subsegment is discarded, otherwise,
5 the subsegment is retained as part of the segment. The procedure 300 then returns to block 314 to repeat the bottom-down filtering steps.

From terminal Z, the procedure 300 may exit to return to an overarching method (such as method 200 described above), or the procedure 300 may return to block 302 to further filter the filtered segments by iterating over the procedure 300 again. This
10 iteration may continue until the procedure 300 does not change the filtered segment.

UTILIZING PRECURSOR PATTERN DETECTION MODELS

Once a precursor pattern detection model has been created, that model can be used while monitoring a patient 90 to detect occurrences of the precursor pattern and
15 cause the associated medical condition to be addressed. FIGURE, 4 is a flowchart that illustrates an exemplary embodiment of a method of using a precursor pattern detection model to detect medical events according to various aspects of the present disclosure. From a start block, the method 400 proceeds to block 402, where the patient monitoring system 102 retrieves one or more precursor pattern detection models from the precursor
20 pattern detection system 104. In some embodiments, the patient monitoring system 102 may be configured to monitor the patient 90 for occurrences of a single medical event, while in other embodiments, the patient monitoring system 102 may be configured to gather as much information as is practical from the patient 90 and to search for precursor patterns indicating multiple medical events.

25 The method 400 proceeds to block 404, where the patient monitoring system 102 collects new data from a patient 90 using one or more sensor devices, including the physical data sensor devices 108 and the interactive sensor devices 110. In some embodiments, fewer sensor devices may be used than during the data collection process, as the precursor pattern detection models may indicate that only some of the collected
30 data was relevant. In some embodiments, more sensor devices than are relevant to the precursor pattern detection models may be used.

At block 406, the patient monitoring system 102 analyzes the new collected data using the one or more precursor pattern detection models. At block 408, in response to

detecting a match between the new collected data and a precursor pattern detection model, the patient monitoring system 102 causes an action to be taken to address the relevant medical event. In some embodiments, the action may include executing the instructions associated with the precursor pattern detection model in the precursor pattern data store 116. The instructions may cause a computing device to inform the patient that they should seek care; may cause a message to be sent to a care provider, may cause a medication to be automatically administered to the patient 90; may cause an amount of monitoring of the patient to be changed, or may cause any other suitable action to occur. As illustrated, the method 400 then proceeds to an end block and terminates, though in some embodiments, the method 400 may run continuously to collect data from patients, analyze the data using the precursor pattern detection models, and take action when appropriate, until the method 400 is terminated by the **patient** 90 or other user.

TEST RESULTS

Multiple tests were performed to verify that the above-described systems and methods can be used to generate precursor pattern detection models for multiple different types of medical events. For each different type of medical event, the above-described method was able to generate effective precursor pattern detection models.

For the first three tests, datasets from the publicly available PhysioNet database (described in Goldberger et al., "Physiobank, physiotoolkit, and physionet: Components of a new research resource for complex physiologic signals," *Circulation*, Volume 101, No. 23, pp. e215-e220 (2000)) were used to build precursor pattern detection models for three different medical events: epileptic seizures, apnea, and vital sign alerts.

The *chbmit* dataset from PhysioNet was used to build the precursor pattern detection model for epileptic seizures. This dataset was collected at the Children's Hospital Boston. It consists of EEG recordings of 22 pediatric subjects with epileptic seizure. The EEG signals are sampled at 256Hz with 16-bit resolution. During the 800 hours of recordings, there are 129 instances of annotated seizure attacks.

The *apnea-ecg* dataset from PhysioNet was used to build the precursor pattern detection model for apnea. This dataset comprises 70 records of a continuous electrocardiography (ECG) signal, sampled at 100Hz with 16-bit resolution, and a set of apnea annotation derived by human experts at 1-minute intervals. The total recording lasts about 500 hours.

The *MIMIC II* dataset was used to build the precursor pattern detection model for vital sign alerts. The dataset is made up of 4448 records from ICU patients. The records include ECG, blood pressure, respiration, and vital signs. The alerts are annotated automatically by an ICU monitor.

5 For each of these datasets, the method 200 discussed above was performed to identify the precursor patterns, and seven well-known classifiers were used to generate the precursor pattern detection models and to validate their performance. The classifiers used were Naive Bayes, Bayes Network, Logistic Regression, C4.5 Decision Tree, Support Vector Machine (SVM), Voting Feature Interval (VFI), and Artificial Neural
10 Network (ANN). Though use of each of these classifiers is well known to one of ordinary skill in the art, brief descriptions of some these classifiers are available in FIGURE 5.

The first metric used to evaluate the performance of the generated precursor pattern detection models was prediction accuracy. Given the limited number of positive
15 segments in the datasets, the experiments were conducted using 10-fold validation. A 10-minute prediction horizon (the initial length of the positive and negative segments) was used for each experiment and the same number of positive and negative segments were used in each case to prevent screwing.

FIGURE 6 is a chart that illustrates the prediction accuracy obtained for these
20 three datasets for each classifier. From these results, it is clear that the prediction accuracy is fairly consistent for all three datasets regardless of the type of classifier used. The highest accuracy of 84.7% is achieved using SVM based on the precursor pattern detection model built for *apnea-ecg*. Overall, C4.5, SVM, and ANN perform slightly better than other classifiers in terms of prediction accuracy. Note that while the results
25 here are comparable to some reported in other studies, the method and system described above do not require any medical domain expertise or analysis to attain this level of accuracy.

The second part of the experiments investigate the false-positive rate, measured in the number of false-positive indications per hour. To measure false-positive rate, the
30 precursor patterns detection models are used to classify all unseen negative segments in the dataset. Ideally, none of these segments should trigger a positive prediction, thus resulting in a zero false-positive rate.

FIGURE 7 is a chart that illustrates the false-positive rate of all three datasets using each different type of classifier. The *chbmit* and *apnea-ecg* datasets produce similar false-positive rates ranging from 0.29/hr to 0.91/hr, which is considerably higher than that of *MIMIC II*. One possible explanation for the difference is misalignment. The annotations in *chbmit* and *apnea-ecg* are done manually, while the *MIMIC II* dataset contains automatic annotations generated by machines. The false-positive rate should be reduced if the annotations are more precisely aligned with the medical events being annotated. In terms of differences between classifiers, C4.5 once again produces the best overall results, followed closely by SVM and ANN. While Logistic Regression achieves the lowest false-positive rate of 0.04/hr for *MIMIC II*, it performs poorly for *chbmit*. This inconsistency suggests that Logistic Regression may not be a suitable classifier for the generalized method.

The tests also showed the ability of the method to discover non-trivial patterns that are not known to experts. For example, when selecting the most prominent features in the *chbmit* dataset, it was discovered that the top-100 features consist entirely of signals acquired from only two EEG channels: F4-C4 and F_{p1}-F₃. In other words, the method was able to determine that these are the only two channels that are relevant for seizure prediction. Likewise, the temporal-spectral distribution of the top-100 most relevant features for the *MIMIC II* dataset indicated that the majority of the most relevant features were concentrated with frequency less than 50Hz and within 5 minutes prior to the alerts. This suggests that a reliable prediction can still be made even if the signals were sampled at a lower sampling frequency and a shorter prediction horizon.

A fourth study was performed to create a precursor pattern detection model for predicting the worsening of heart failure symptoms. Previously, the most widely used metric to predict the worsening of heart failure symptoms was daily weight change (DWC). It is even recommended by the American College of Cardiology/American Health Association as a potential indicator for water retention, which leads to swollen ankles and other heart failure symptoms. However, several studies have shown that DWC has relatively weak correlation with the worsening of heart failure symptoms. Clinical trials of remote patient monitoring symptoms have also anecdotally shown this poor predictive power: when medical staff followed up with patients upon observing a DWC of more than two pounds, the patient almost always attributed the change to food or normal weight fluctuation, and denied any worsening of heart failure symptoms.

The case for using DWC as a predictor for the worsening of heart failure symptoms is further weakened from the charts illustrated in FIGURES 8A and 8B. These figures show systolic and diastolic blood pressure (SBP & DBF) and weight readings from the same patient over two separate 7-day periods. By the seventh day in
5 FIGURE 8A, the patient is known to have reported worsening of heart failure symptoms, whereas the patient reported no change in symptoms during the seven day period in FIGURE 8B. Although the weight fluctuated in both cases, there is very limited change in weight (< lib) measured on the seventh day. If DWC were used to predict the symptom, there would most probably end up having either a large number of false
10 positives, or the true positives would be missed altogether. On the other hand, the charts suggest that the systolic and diastolic blood pressures are less stable before the worsening of heart symptoms, suggesting that these values may serve as better predictors for this medical event.

Accordingly, for the fourth study, remote patient monitoring data was collected
15 for study participants experiencing heart failure. The population of participants was 89% male with a mean age of approximately 54. Study participants were provided with Bluetooth weight scales, blood pressure monitors, heart rate monitors, landline or Ethernet gateways, and Android activity monitoring applications. Patients also answer a series of nine questions on a daily basis, the answers to which are interpreted to
20 determine whether symptoms have stabilized or have worsened. These answers were crosschecked with nurses' call logs to remove any accidental false reporting. Based on these criteria there were a total of 34 instances of worsening of heart failure symptoms self-reported by nine patients during the clinical trials.

These 34 instances of worsening of heart failure symptoms were annotated in the
25 collected data, and used to extract the positive segments. All other segments that are at least three days away from the instances were extracted as negative segments. Instances that occurred within three days of other instances were discarded. A total of nine features were extracted for each segment: daily change in systolic blood pressure, daily change in diastolic blood pressure, daily change in weight, standard deviation of systolic blood
30 pressure over the past three days, standard deviation of diastolic blood pressure over the past three days, standard deviation of weight over the past three days, standard deviation of systolic blood pressure over the past seven days, standard deviation of diastolic blood pressure over the past seven days, and standard deviation of weight over the past seven

days. Six different classifiers, listed in FIGURE 5, were used. These classifiers were chosen because they utilize a wide range of fundamentally different statistical principles. This ensures that the results are based on the intrinsic discriminatory power of the extracted features rather than a particular algorithm overfitting the model. To compare, the traditional method of detecting worsening of heart failure symptoms upon detecting a
5 DWC of two pounds was also used.

Several measures were taken to ensure the fairness of the experiments. Firstly, equal numbers of positive and negative instances were used in each experiment to prevent biasing toward a particular class. Secondly, the results were computed using ten-fold
10 cross validation to guard against overfitting given the relatively small number of instances. Finally, each experiment was conducted ten times and the average results are reported, thus reducing the effect of outliers.

The algorithms were evaluated in terms of their accuracy, sensitivity, and specificity. They are computed based on the number of True Positives (TP), True
15 Negatives (TN), False Positives (FP), and False Negatives (FN) in the prediction results. The results for each algorithm are plotted in Figure 4. Since the number of positive and negative instances is balanced, the minimal accuracy of any better-than-chance algorithm should be greater than 0.5. The chart in FIGURE 9 shows that predicting the worsening of heart failure symptoms using DWC seems to be only marginally more accurate than
20 guessing, with an accuracy of 0.519. This is not surprising considering that only 15.9% of the positive instances were classified as such, whereas a substantial amount of negative instances (28.6%) turn out to have a DWC of more than 2 lbs.

On the other hand, the precursor pattern detection models calculated as described above attained much higher accuracy. Three algorithms, NBC, LR, and RIDOR, were
25 able to correctly predict the worsening or stabilization of heart failure symptoms 74% of the time. Even though VFI has the lowest accuracy of 0.696, it was still nearly 20% more accurate than DWC. Furthermore, all six algorithms had sensitivity values that were at least 45% greater than that of DWC. This suggests that the extracted features have a much stronger correlation to the worsening of heart failure symptoms than simple daily
30 weight change, which can be further confirmed by the high specificity ranged from 0.696 (for kNN) to 0.87 (for RIDOR).

However, not all the features were equally predictive. Out of the nine features, the study determined that the classifiers consistently identify three- and seven-day

standard deviations of systolic and diastolic blood pressure as the most predictive ones. Actually, C4.5, RiDOR, and LR even ignore the other features entirely when constructing the models.

FIGURE 10 is a chart that illustrates Receiver Operating Characteristic (ROC) curves for the classifiers and DWC. These curves demonstrate the effect of varying the classification algorithm's discrimination threshold on the true positive rate (TPR) vs. the false positive rate (FPR). A perfect algorithm will produce a point at the top left corner of the figure where $TPR = 1$ and $FPR = 0$. On the other hand, an algorithm that guesses randomly will produce a curve that follows the dotted 45-degree line in the figure where $TPR = FPR$. Therefore, the area under the curve (AUC) is often used as a metric to gauge the performance of a classification algorithm.

It is dear from the chart in FIGURE 10 that DWC is not significantly better than random guessing. In fact, the AUC for DWC is estimated to be 0.573, which is less than 8% higher than that of a completely random guess. In comparison, the other classifiers used perform significantly better. The AUC ranged from 0.817 for LR to 0.621 for VFI. It is worth noting that NBC can be tuned to achieve a decent TPR of 0.652 while minimizing FPR to 0 at the same time. This makes NBC the most suitable algorithm to use for this particular application when the consequence of a false alarm is more costly than that of a missed prediction.

20

EXEMPLARY COMPUTING DEVICE

FIGURE 11 is a block diagram that illustrates aspects of an exemplary computing device 1100 appropriate for use with embodiments of the present disclosure. While FIGURE 11 is described with reference to a computing device that is implemented as a device on a network, the description below is applicable to servers, personal computers, mobile phones, smart phones, tablet computers, embedded computing devices, and other devices that may be used to implement portions of embodiments of the present disclosure. Moreover, those of ordinary skill in the art and others will recognize that the computing device 1100 may be any one of any number of currently available or yet to be developed devices.

30

In its most basic configuration, the computing device 1100 includes at least one processor 1102 and a system memory 1104 connected by a communication bus 1106. Depending on the exact configuration and type of device, the system memory 1104 may

be volatile or nonvolatile memory, such as read only memory ("ROM"), random access memory ("RAM"), EEPROM, flash memory, or similar memory technology. Those of ordinary skill in the art and others will recognize that system memory 1104 typically stores data and/or program modules that are immediately accessible to and/or currently
5 being operated on by the processor 1102. In this regard, the processor 1102 may serve as a computational center of the computing device 1100 by supporting the execution of instructions.

As further illustrated in FIGURE 11, the computing device 1100 may include a network interface 1110 comprising one or more components for communicating with
10 other devices over a network. Embodiments of the present disclosure may access basic sendees that utilize the network interface 1110 to perform communications using common network protocols. The network interface 1110 may also include a wireless network interface configured to communicate via one or more wireless communication protocols, such as WiFi, 2G, 3G, LTE, WiMAX, Bluetooth, and/or the like.

In the exemplary embodiment depicted in FIGURE 11, the computing device 1100 also includes a storage medium 1108. However, services may be accessed using a computing device that does not include means for persisting data to a local storage medium. Therefore, the storage medium 1108 depicted in FIGURE 11 is
15 represented with a dashed line to indicate that the storage medium 1108 is optional. In any event, the storage medium 1108 may be volatile or nonvolatile, removable or nonremovable, implemented using any technology capable of storing information such as, but not limited to, a hard drive, solid state drive, CD ROM, DVD, or other disk storage, magnetic cassettes, magnetic tape, magnetic disk storage, and/or the like.
20

As used herein, the term "computer-readable medium" includes volatile and non-volatile and removable and non-removable media implemented in any method or
25 technology capable of storing information, such as computer readable instructions, data structures, program modules, or other data. In this regard, the system memory 1104 and storage medium 1108 depicted in FIGURE 11 are merely examples of computer-readable media.

Suitable implementations of computing devices that include a processor 1102, system memory 1104, communication bus 1106, storage medium 1108, and network
30 interface 1110 are known and commercially available. For ease of illustration and because it is not important for an understanding of the claimed subject matter,

FIGURE 11 does not show some of the typical components of many computing devices. In this regard, the computing device 1100 may include input devices, such as a keyboard, keypad, mouse, microphone, touch input device, touch screen, tablet, and/or the like. Such input devices may be coupled to the computing device 1100 by wired or wireless connections including RF, infrared, serial, parallel, Bluetooth, USB, or other suitable connections protocols using wireless or physical connections. Similarly, the computing device 1100 may also include output devices such as a display, speakers, printer, etc. Since these devices are well known in the art, they are not illustrated or described further herein.

10 As will be appreciated by one skilled in the art, the specific routines described above in the flowcharts may represent one or more of any number of processing strategies such as event-driven, interrupt-driven, multi-tasking, multi-threading, and the like. As such, various acts or functions illustrated may be performed in the sequence illustrated, in parallel, or in some cases omitted. Likewise, the order of processing is not necessarily required to achieve the features and advantages, but is provided for ease of illustration and description. Although not explicitly illustrated, one or more of the illustrated acts or functions may be repeatedly performed depending on the particular strategy being used. Further, these FIGURES may graphically represent code to be programmed into a computer readable storage medium associated with a computing device.

20 While illustrative embodiments have been illustrated and described, it will be appreciated that various changes can be made therein without departing from the spirit and scope of the claimed subject matter.

CLAIMS

The embodiments of the invention in which an exclusive property or privilege is claimed are defined as follows:

1. A computer-implemented method of generating a precursor pattern detection model for detecting an occurrence of a given type of medical event, the method comprising:

obtaining annotated monitoring data by a computing device, wherein the annotated monitoring data represents measurements of physical conditions of at least one patient;

extracting, by a computing device, a set of positive segments and a set of negative segments from the annotated monitoring data, wherein each positive segment is associated with an occurrence of a medical event of the given type, and wherein each negative segment is not associated with an occurrence of the medical event of the given type;

executing, by a computing device, a time filter on the set of positive segments to generate a set of filtered positive segments;

executing, by a computing device, a time filter on the set of negative segments to generate a set of filtered negative segments;

extracting, by a computing device, a set of extracted positive features from the set of filtered positive segments and a set of extracted negative features from the set of filtered negative segments;

selecting, by a computing device, a set of selected positive features from the set of extracted positive features and a set of selected negative features from the set of extracted negative features;

calculating, by a computing device, a precursor pattern detection model based on the set of selected positive features and the set of selected negative features; and

storing, by a computing device on a nontransitory computer-readable medium, the precursor pattern detection model.

2. The computer-implemented method of Claim 1, wherein the annotated monitoring data includes annotations indicating occurrences of the given medical event during collection of the monitoring data.

3. The computer-implemented method of Claim 1, wherein the set of selected positive features includes features extracted from less than all of a set of signals included in the annotated monitoring data.

4. The computer-implemented method of Claim 1, wherein executing a time filter on the set of positive segments to generate a set of filtered positive segments includes:

executing a top-down time filter on the set of positive segments to generate a set of top-down filtered positive segments; and

executing a bottom-up time filter on the set of top-down filtered positive segments to generate the set of filtered positive segments.

5. The computer-implemented method of Claim 4, wherein executing a time filter on the set of positive segments to generate a set of filtered positive segments includes:

repeatedly executing the top-down time filter and the bottom-up time filter until a strength of the filtered positive segments does not increase farther.

6. The computer-implemented method of Claim 1, wherein extracting a set of extracted positive features from the set of filtered positive segments includes performing a wavelet transform for each filtered positive segment of the set of filtered positive segments.

7. The computer-implemented method of Claim 6, wherein the wavelet transform is a Haar wavelet transform.

8. The computer-implemented method of Claim 1, wherein extracting a set of extracted positive features from the set of filtered positive segments includes determining a change in a value of each filtered positive segment from a value of a previous positive segment.

9. The computer-implemented method of Claim 1, wherein extracting a set of extracted positive features from the set of filtered positive segments includes determining a standard deviation for each filtered positive segment.

10. The computer-implemented method of Claim 1, wherein selecting a set of selected positive features from the set of extracted positive features includes selecting positive features using a wrapper algorithm.

11. The computer-implemented method of Claim 10, wherein selecting positive features using a wrapper algorithm includes calculating an information gain for each extracted positive feature of the set of extracted positive features.

12. The computer-implemented method of Claim 11, wherein selecting positive features using a wrapper algorithm further includes selecting extracted positive features having an information gain greater than a predetermined threshold value.

13. The computer-implemented method of Claim 1, wherein calculating a precursor pattern detection model based on the set of selected positive features and the set of selected negative features includes providing the set of selected positive features and the set of selected negative features as input to a classifier.

14. The computer-implemented method of Claim 13, wherein the classifier is selected from a group consisting of a naive classifier, a Bayes network classifier, a logistic regression classifier, a C4.5 decision tree classifier, an support vector machine (SVM) classifier, a voting feature interval (VFI) classifier, and an artificial neural network (ANN) classifier.

15. A nontransitory computer-readable medium having computer-executable instructions stored thereon that, in response to execution by one or more processors of one or more computing devices, cause the one or more computing devices to perform the method recited in any of Claims 1-14.

16. A system comprising:

one or more computing devices, each of the one or more computing devices including at least one processor and a memory;

wherein the one or more computing devices are configured to perform the method recited in any of Claims 1-14.

17. A method of detecting an occurrence of a given type of medical event, the method comprising:

retrieving a precursor pattern detection model from a precursor pattern data store;
using one or more monitoring devices to collect data from a patient;
analyzing the collected data using the precursor pattern detection model; and
in response to the analysis indicating that a precursor pattern has been detected in
the collected data, causing an alert to be presented to a caregiver.

18. The method of Claim 17, wherein the precursor pattern detection model is generated using the method recited in any of Claims 1-14.

19. The method of Claim 17, wherein the one or more monitoring devices include at least one monitoring device selected from a group consisting of a sphygmomanometer, a thermometer, a heart rate monitor, a pedometer, an electronic scale, an electroencephalograph machine, an electrocardiogram machine, a motion sensor, and an interactive computing device.

20. The method of Claim 17, wherein the given type of medical event is selected from a group consisting of a seizure, a sleep apnea event, cardiac decompensation, and an arrhythmia.

21. A computing device comprising at least one processor and a memory, wherein the computing device is configured to perform the method recited in any of Claims 17-20.

22. A nontransitory computer-readable medium having computer-executable instructions stored thereon that, in response to execution by one or more processors of a computing device, cause the computing device to perform the method recited in any of Claims 17-20.




Fig. 1.




Fig. 2A.

3/12

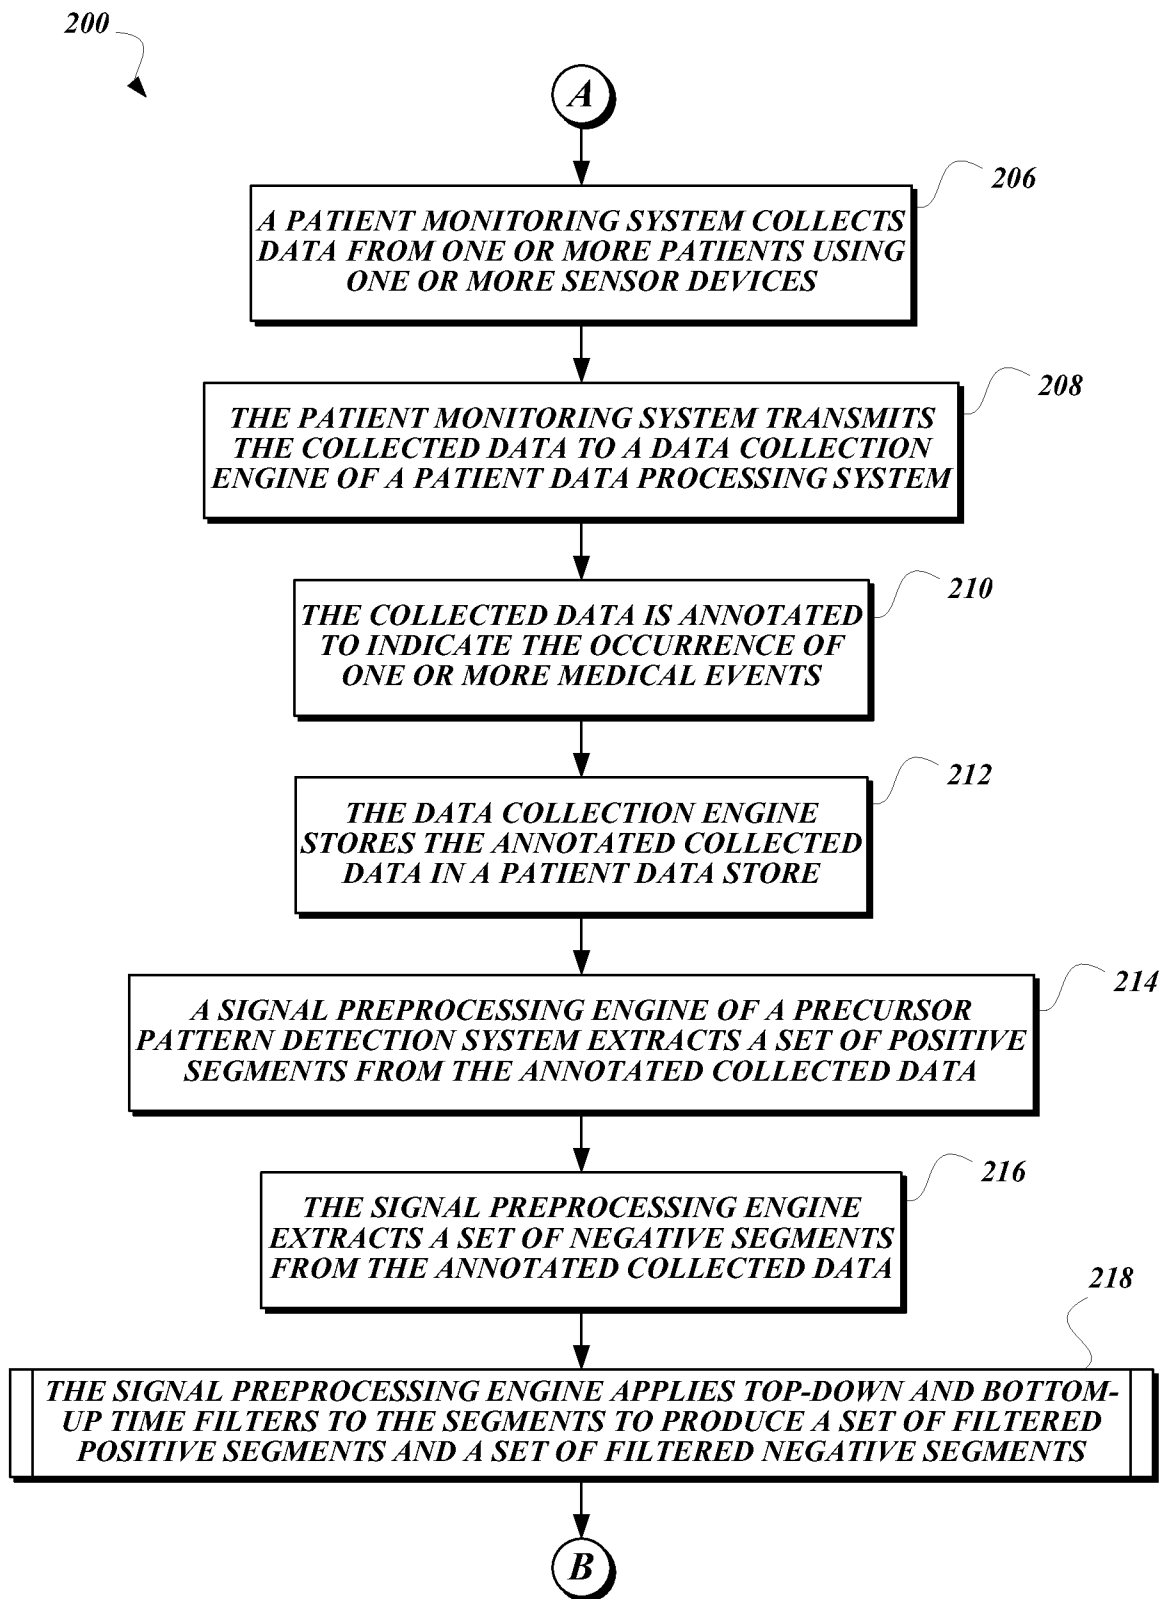


Fig. 2B.

4/12

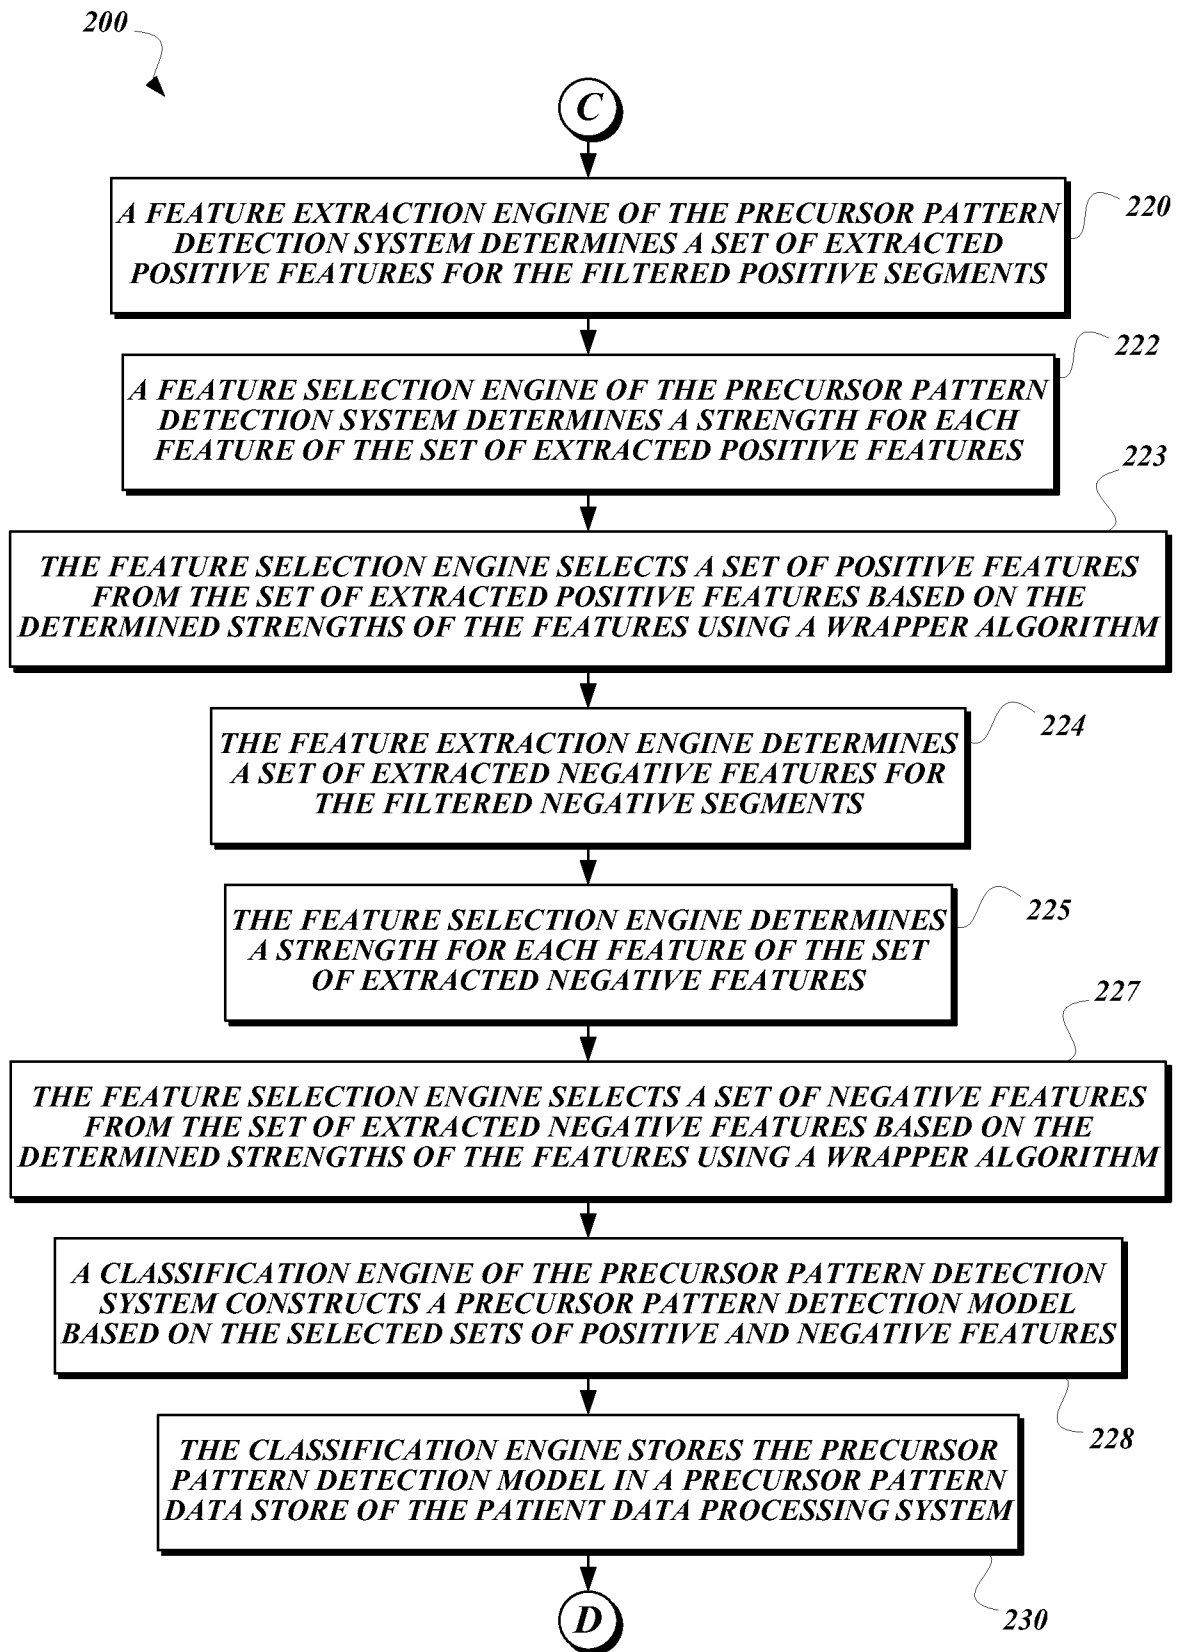


Fig.2C.

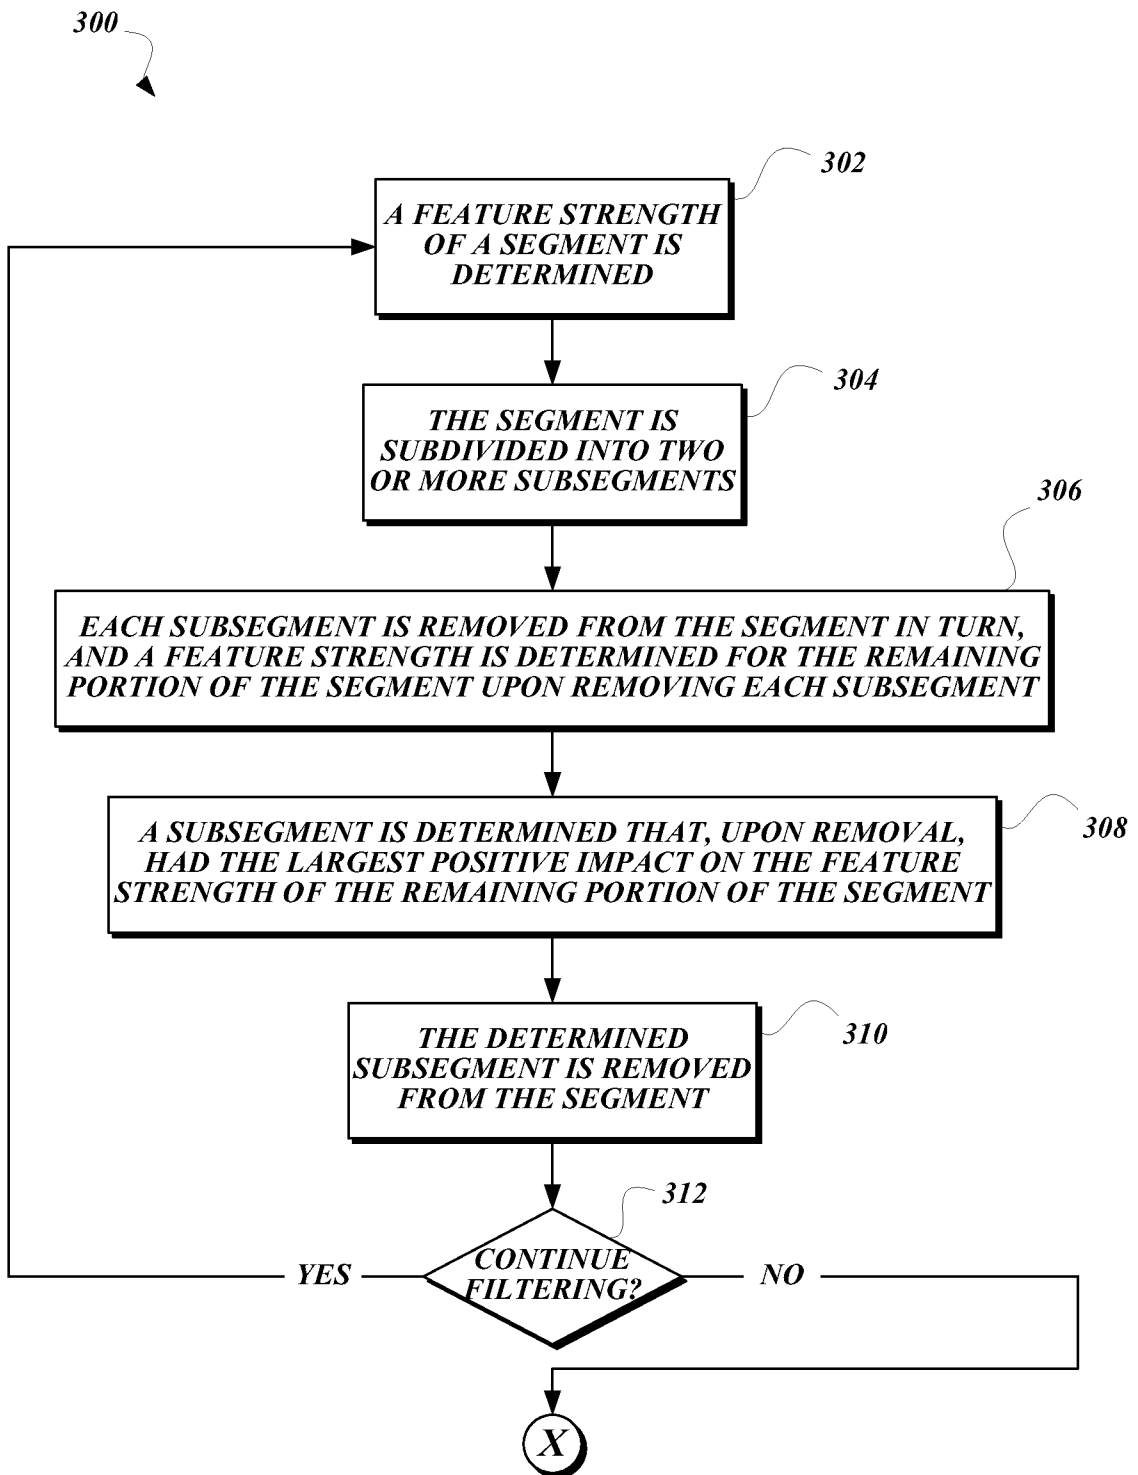


Fig.3A.

6/12

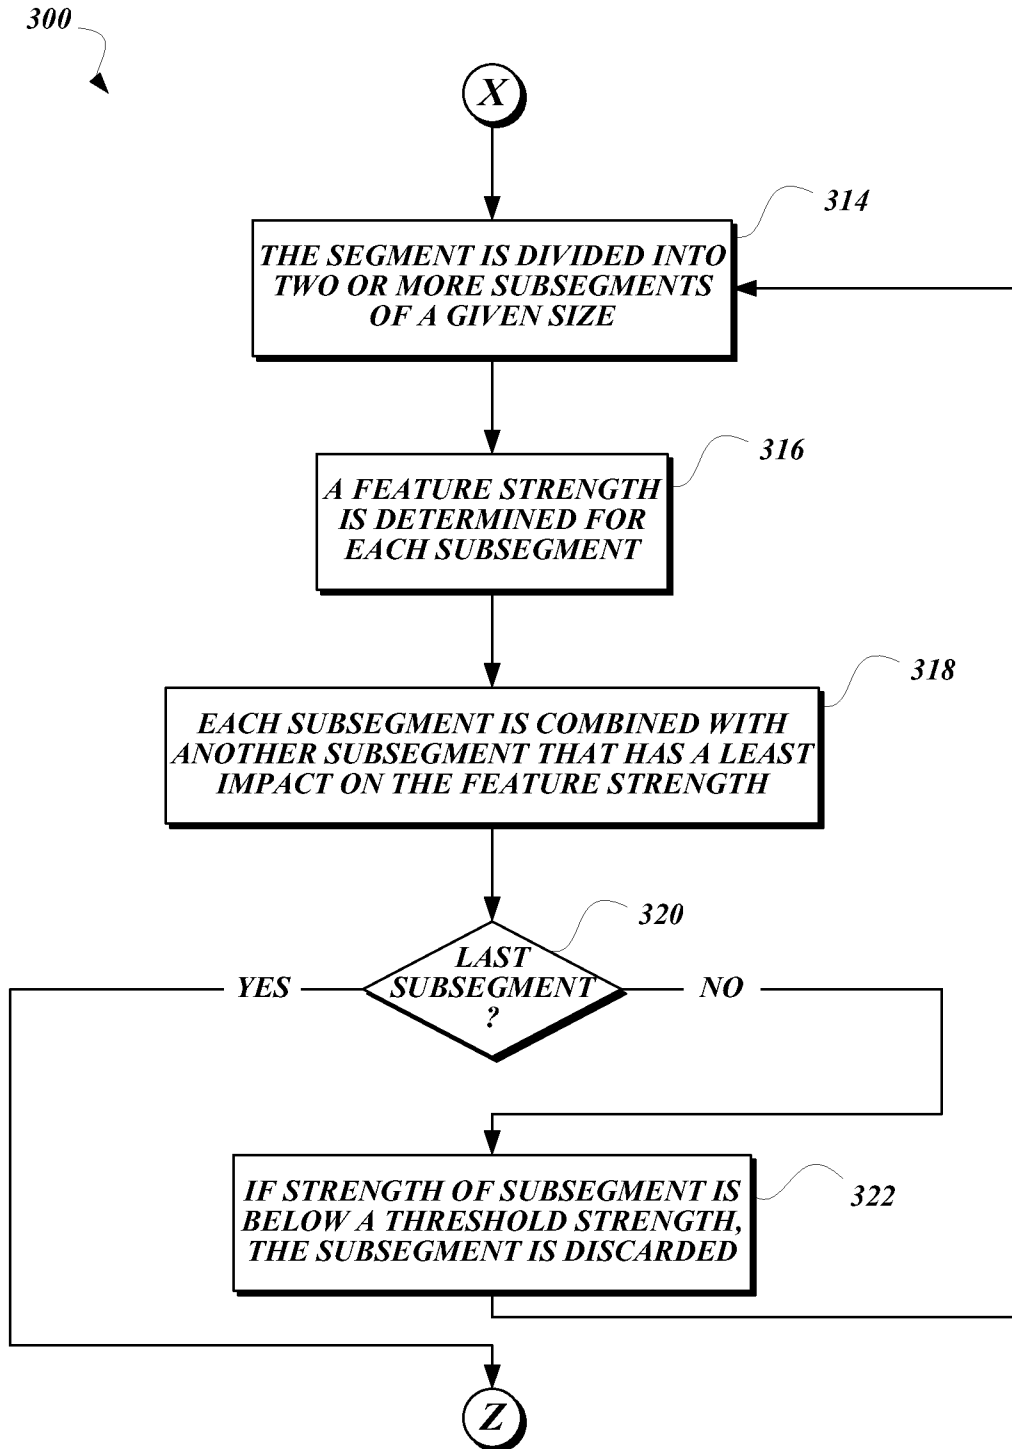


Fig.3B.

7/12

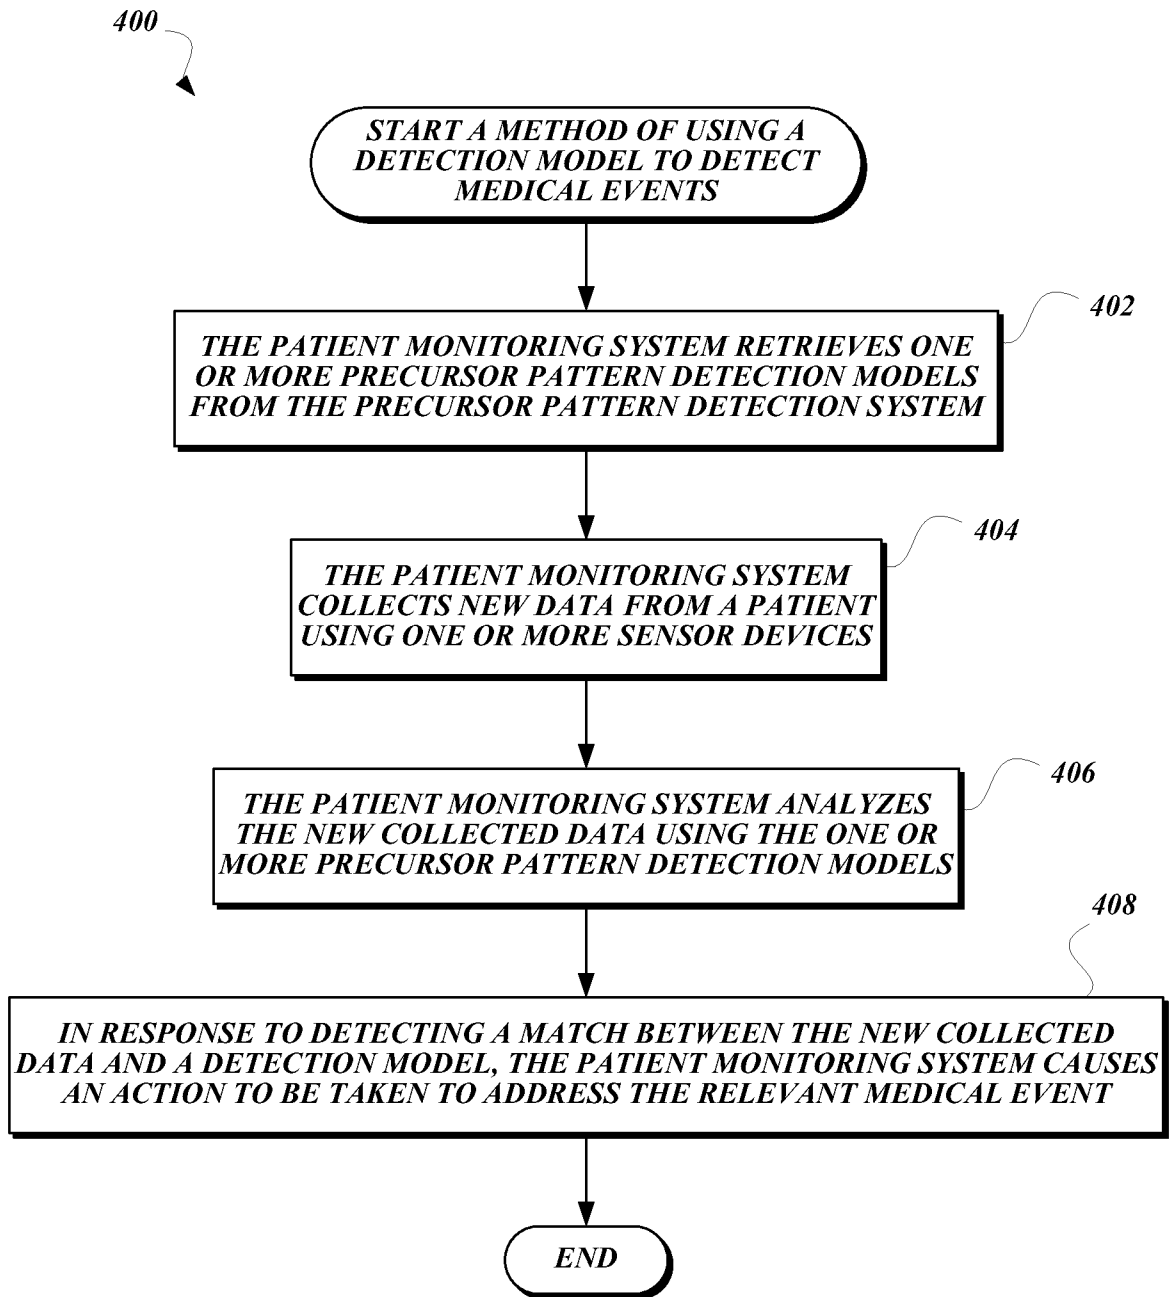


Fig.4.

Algorithm	Description	Parameters
Naive Bayes Classifier (NBC)	<p>NBC uses Bayes' theorem and assumes each feature is conditionally independent. The posterior probability of a class C given a set of features F1, F2, ... Fn is therefore</p> $p(C F_1, F_2, \dots, F_n) \propto p(C) \prod_{i=1}^n p(F_i C)$ <p>The prior and likelihood can be calculated from the training set directly and applies to the test set.</p>	
Nearest Neighbor (kNN)	<p>An unlabeled instance is classified based on its nearest k neighbors from the training set based on majority vote. Euclidean distance of the features is often used to determine the closeness of two instances</p> $d = \sqrt{\sum_{i=1}^n (x_i - y_i)^2}$ <p>where x_i and y_i are corresponding features from the two instances.</p> <p>Logistic regression is a type of regression that predicts the outcome of a binary depend variable (class) based on a set of independent variables (features) X_i using the logistic function</p> $p(x) = \frac{1}{1 + e^{-f(x)}}, \quad f(x) = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \dots + \beta_n X_n$ <p>where $p(x)$ is the probability of x being class 1, given regression coefficients $\beta_1 \dots \beta_n$ and intercept β_0.</p> <p>VFI classifies builds upper and lower bounds around each class for each feature. Classification is based on majority voting, where a vote for class C based from feature a is computed as</p> $v(a, C) = \frac{\text{interval_class_count}(a, i, C)}{\text{class_count}(C)} \left(\frac{H(C a)}{\text{max uncertainty}} \right)^{\text{bias}}$	k = 5
Logistic Regression (LR)		
Voting Feature Interval (VFI)		bias = 0.6
Ripple-Down Rule Learner (RIDOR)	<p>RIDOR is a version of Ripple Down Rule using the Induct algorithm where the default rules are first generated based on least error rate. A set of exception rules are generated to predict classes other than those covered by the default rules.</p> <p>C4.5 constructs a decision tree based on information entropy. Each feature in the training set is evaluated for its information gain</p> $IG(T, a) = H(T) - H(T a) \quad H(T) = -\sum_{i=1}^n p(x_i) \ln p(x_i)$ <p>Feature with the largest IG and have yet been used is chosen to split the decision tree at each node until the confidence falls below certain threshold.</p>	confidence = 0.25
C4.5 Decision Tree (C4.5)		

Fig.5.

9/12

	<i>CHBMIT</i>	<i>APNEA-ECG</i>	<i>MIMIC II</i>
<i>NAÏVE</i>	77.8%	81.5%	78.7%
<i>BAYES NETWORK</i>	79.8%	80.6%	79.3%
<i>LOGISTIC REGRESSION</i>	76.7%	79.3%	75.4%
<i>C4.5</i>	81.3%	82.1%	80.9%
<i>SVM</i>	80.2%	84.7%	82.1%
<i>VFI</i>	79.4%	79.7%	78.8%
<i>ANN</i>	79.8%	81.5%	81.3%

Fig. 6.

	<i>CHBMIT</i>	<i>APNEA-ECG</i>	<i>MIMIC II</i>
<i>NAÏVE</i>	0.52/HR	0.41/HR	0.17/HR
<i>BAYES NETWORK</i>	0.49/HR	0.39/HR	0.11/HR
<i>LOGISTIC REGRESSION</i>	0.91/HR	0.33/HR	0.04/HR
<i>C4.5</i>	0.33/HR	0.29/HR	0.05/HR
<i>SVM</i>	0.37/HR	0.27/HR	0.13/HR
<i>VFI</i>	0.45/HR	0.38/HR	0.20/HR
<i>ANN</i>	0.51/HR	0.44/HR	0.06/HR

Fig. 7.

10/12




Fig.8A.




Fig.8B.

11/12




Fig.9.




Fig.10.




Fig.11.

A. CLASSIFICATION OF SUBJECT MATTER**G06F 19/00(2011.01)i**

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 19/00; A61B 5/04; G06Q 50/22; G06Q 50/00; A01B 5/00

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

eKOMPASS(KIPO internal) & Keywords precursor pattern detection, medical event, given type, patient, negative segment, annotated monitoring data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	KR 10-2010-0126365 A (PHADIA AB) 01 December 2010 See paragraphs [0067]- [0137] ; claim V, and figure 1.	1-16
A	US 6549804 B1 (IVAN OSORIO et al.) 15 April 2003 See column 8, line 40 - column 10, line 2; claim V, and figure 1.	1-16
A	JP 2007-534086 A (BIO- RAD LAB INC.) 22 November 2007 See abstract ; paragraphs [0009]- [0012] ; and figure 1.	1-16
A	JP 2007-066058 A (TOSHIBA CORP.) 15 March 2007 See paragraphs [0010]- [0020] ; claim V, and figures 1-8 .	1-16
A	US 6248063 B1 (STEPHEN D. BARNHILL et al.) 19 June 2001 See abstract ; column 11, line 14 - column 13, line 60; and figures 1-4 .	1-16

I 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

"E" earlier application or patent but published on or after the international filing date

"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)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"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

"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

"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

"&" document member of the same patent family


Date of the actual completion of the international search

24 December 2013 (24.12.2013)

Date of mailing of the international search report

27 December 2013 (27.12.2013)

Name and mailing address of the ISA/KR



 Korean Intellectual Property Office
 189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan City,
 302-701, Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

HWANG, Yun Koo

Telephone No. +82-42-481-5715



Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. Claims Nos. : 17-20
because they relate to subject matter not required to be searched by this Authority, namely:
Claims 17-20 pertain to methods for treatment of human body by therapy, as well as diagnostic methods, and thus relate to a subject matter which this International Searching Authority is not required to search under Article 17(2)(a)(i) and Rule 39.1(iv) PCT.
2. Claims Nos. :
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically :
3. Claims Nos. : 21-22
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of any additional fees.
3. As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos. :

Remark on Protest

- The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/057137

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
KR 10-2010-0126365 A	01/12/2010	AU 2009-210848 AI	13/08/2009
		CA 2713602 AI	13/08/2009
		CN 101971180 A	09/02/2011
		EP 2255308 AI	01/12/2010
		JP 2011-511297 A	07/04/2011
		US 2010-0332143 AI	30/12/2010
		us 8437963 B2	07/05/2013
		wo 2009-099379 AI	13/08/2009
US 6549804 BI	15/04/2003	AU 1752897 A	20/08/1997
		DE 69736592 DI	12/10/2006
		DE 69736592 T2	14/12/2006
		EP 0898460 AI	03/03/1999
		EP 0898460 A4	24/03/1999
		EP 0898460 BI	30/08/2006
		JP 2002-502270 A	22/01/2002
		JP 3769023 B2	19/04/2006
		US 2005-0197590 AI	08/09/2005
		us 2010-0198098 AI	05/08/2010
		us 5995868 A	30/11/1999
		us 7630757 B2	08/12/2009
wo 97-26823 AI	31/07/1997		
JP 2007-534086 A	22/11/2007	AU 2001-296883 A8	29/04/2002
		AU 2001-96883 AI	29/04/2002
		AU 2005-236067 AI	03/11/2005
		CA 2424764 AI	25/04/2002
		CA 2562530 AI	03/11/2005
		CN 1961212 A	09/05/2007
		EP 1327148 A2	16/07/2003
		EP 1327148 BI	15/04/2009
		EP 1738168 A2	03/01/2007
		EP 1738168 A4	10/06/2009
		JP 2004-511805 A	15/04/2004
		US 2004-0265874 AI	30/12/2004
		wo 02-033415 A3	27/03/2003
		wo 02-33415 A2	25/04/2002
		wo 2005-103300 A2	03/11/2005
		wo 2005-103300 A3	20/07/2006
wo 2005-103300 A9	11/05/2006		
JP 2007-066058 A	15/03/2007	JP 4381361 B2	09/12/2009
		us 2007-0079195 AI	05/04/2007
		us 7437266 B2	14/10/2008
US 6248063 BI	19/06/2001	AU 1837495 A	06/05/1996
		AU 1996-67131 B2	23/03/2000
		CA 2227543 AI	13/02/1997
		CA 2227543 C	05/10/2010

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/057137

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		CN 1194045 AO	23/09/1998
		EP 0842475 AI	14/04/1999
		EP 0842475 BI	08/11/2000
		JP 11-504739 A	27/04/1999
		JP 2007-052774 A	01/03/2007
		US 5769074 A	23/06/1998
		us 6306087 BI	23/10/2001
		wo 96-12187 AI	25/04/1996
		wo 97-05553 AI	13/02/1997

专利名称(译)	用于计算和使用统计模型来预测医疗事件的方法和系统		
公开(公告)号	EP2891094A1	公开(公告)日	2015-07-08
申请号	EP2013832651	申请日	2013-08-28
[标]申请(专利权)人(译)	加利福尼亚大学董事会		
申请(专利权)人(译)	加利福尼亚大学董事会		
当前申请(专利权)人(译)	加利福尼亚大学董事会		
[标]发明人	SARRAFZADEH MAJID SUH MYUNG KYUNG LAN MARS GHASEMZADEH HASSAN		
发明人	SARRAFZADEH, MAJID SUH, MYUNG-KYUNG LAN, MARS GHASEMZADEH, HASSAN		
IPC分类号	G06F19/00 A61B5/00 A61B5/0205 A61B5/0402 A61B5/0476 A61B5/11 G06N99/00		
CPC分类号	A61B5/7282 A61B5/02055 A61B5/0402 A61B5/0476 A61B5/112 A61B5/4094 A61B5/4818 A61B5/725 A61B5/726 A61B5/7264 A61B5/746 F04C2270/041 G06N20/00 G16H50/20		
代理机构(译)	HOARTON , LLOYD DOUGLAS CHARLES		
优先权	61/694171 2012-08-28 US		
其他公开文献	EP2891094A4		
外部链接	Espacenet		

摘要(译)

公开了用于广泛前体模式发现的系统和方法，其与广泛的生物医学信号和应用一起工作以检测广泛的医学事件。在一些实施例中，所述方法和系统不需要基于所分析的医疗事件的特定领域知识或重要的重新配置，因此还可以发现先前专家未知的模式。在一些实施例中，为了构建前体模式检测模型，系统获得带注释的监测数据。从带注释的监视数据中提取正片段和负片段，并进行预处理。从预处理的片段中提取特征，并从提取的特征中选择所选特征。对所选特征进行分类以创建前体模式检测模型。然后可以实时使用前体模式检测模型来检测感兴趣的医疗事件的发生。