

DROWSINESS ONSET DETECTION

BACKGROUND

[001] Electroencephalography (EEG) is a method of monitoring electrical activity of the brain over time, as captured using multiple electrodes placed on the scalp. Information about a person's mental and physical condition can be derived from an analysis of the so-called brain waves observed in EEG signals. One such condition is the transition between a wakeful state and a drowsy state.

[002] In human sleep, there are two main types of sleep—namely rapid-eye-movement (REM) sleep and non-rapid-eye-movement (NREM) sleep. REM sleep is identifiable in an EEG signal by its characteristic low-amplitude, high-frequency waves. NREM sleep is categorized in three distinct stages: N1, N2, and N3; and presents very different EEG wave patterns from REM sleep. For example, N3 NREM sleep exhibits brain waves in an EEG signal characterized by high-amplitude, low-frequency waves. In general, brain waves become slower and more synchronized as a person transitions from stage N1 to N3 sleep.

[003] As a person falls asleep, the activity of their body slows down and their brain waves get slower and bigger. N1 NREM sleep, which is often referred to as drowsy sleep or drowsiness, usually occurs between sleep and wakefulness. This stage is characterized in EEG signals as a transition between so-called alpha waves having a frequency of 8–13 Hz (which is common in the awake state) to so-called theta waves having a frequency of 4–7 Hz. Thus, the transition from a wakeful state to a drowsy state can be easily seen in an EEG signal.

SUMMARY

[004] The drowsiness onset detection implementations described herein generally detect the onset of drowsiness in an individual. In one general implementation this involves using one or more heart rate (HR) sensors which capture HR information of the individual, one or more computing devices and a drowsiness onset detector computer program having a plurality of sub-programs executable by the one or more computing devices. The one or more computing devices are directed by the sub-programs of the drowsiness onset detector computer program to first receive the HR information from the heart rate sensor or sensors. A set of features is then extracted from the HR information. These features, which are among many features that can be extracted from HR information, have been determined to be specifically indicative of a transition from a wakeful state of an

individual to a drowsy state of an individual. Next, the extracted features are combined to produce a drowsiness detection input. The drowsiness detection input is then fed into an artificial neural network (ANN) classifier that has been trained to distinguish between the wakeful state of an individual and the drowsy state of an individual based on the extracted features. In one implementation, the ANN classifier has also been trained in part using an indicator of the transition between wakefulness and drowsiness (i.e., N1 NREM sleep) as exhibited in an EEG signal. It is next identified from an output of the ANN classifier if the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness. If so, a drowsiness onset warning is initiated.

[005] It should be noted that the foregoing 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 or essential 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. Its sole purpose is to present some concepts of the claimed subject matter in a simplified form as a prelude to the more-detailed description that is presented below.

DESCRIPTION OF THE DRAWINGS

[006] The specific features, aspects, and advantages of the disclosure will become better understood with regard to the following description, appended claims, and accompanying drawings where:

[007] FIG. 1 is a diagram illustrating one implementation, in simplified form, of a system framework for realizing the drowsiness onset detection implementations described herein.

[008] FIG. 2 is a flow diagram illustrating an exemplary implementation, in simplified form, of a process for detecting the onset of drowsiness for an individual.

[009] FIG. 3 is a diagram illustrating another implementation, in simplified form, of a system framework for realizing the drowsiness onset detection implementations described herein, which involves a remote server or cloud service.

[0010] FIG. 4 is a flow diagram illustrating an exemplary implementation, in simplified form, of a process for detecting the onset of drowsiness for an individual, which involves a remote server or cloud service.

[0011] FIG. 5 is a diagram illustrating an exemplary implementation, in simplified form, of a system framework for training a drowsiness onset detection classifier, which in turn can be used to detect the onset of drowsiness in an individual.

[0012] FIG. 6 is a flow diagram illustrating an exemplary implementation, in simplified form, of a process for training a drowsiness onset detection classifier.

[0013] FIG. 7 is a diagram depicting a general purpose computing device constituting an exemplary system for use with the drowsiness onset detection implementations described herein.

DETAILED DESCRIPTION

[0014] In the following description reference is made to the accompanying drawings which form a part hereof, and in which are shown, by way of illustration, specific versions in which drowsiness onset detection implementations can be practiced. It is understood that other implementations can be utilized and structural changes can be made without departing from the scope thereof.

[0015] It is also noted that for the sake of clarity specific terminology will be resorted to in describing the drowsiness onset detection implementations and it is not intended for these implementations to be limited to the specific terms so chosen. Furthermore, it is to be understood that each specific term includes all its technical equivalents that operate in a broadly similar manner to achieve a similar purpose. Reference herein to “one implementation”, or “another implementation”, or an “exemplary implementation”, or an “alternate implementation” means that a particular feature, a particular structure, or particular characteristics described in connection with the implementation can be included in at least one version of the drowsiness onset detection. The appearances of the phrases “in one implementation”, “in another implementation”, “in an exemplary implementation”, and “in an alternate implementation” in various places in the specification are not necessarily all referring to the same implementation, nor are separate or alternative implementations mutually exclusive of other implementations. Yet furthermore, the order of process flow representing one or more implementations of the project information extraction does not inherently indicate any particular order or imply any limitations thereof.

[0016] As utilized herein, the terms “component,” “system,” “client” and the like are intended to refer to a computer-related entity, either hardware, software (e.g., in execution), firmware, or a combination thereof. For example, a component can be a process running on a processor, an object, an executable, a program, a function, a library, a subroutine, a computer, or a combination of software and hardware. By way of illustration, both an application running on a server and the server can be a component. One or more components can reside within a process and a component can be localized on one

computer and/or distributed between two or more computers. The term "processor" is generally understood to refer to a hardware component, such as a processing unit of a computer system.

[0017] Furthermore, to the extent that the terms "includes," "including," "has," "contains," and variants thereof, and other similar words are used in either this detailed description or the claims, these terms are intended to be inclusive in a manner similar to the term "comprising" as an open transition word without precluding any additional or other elements.

1.0 Drowsiness Onset Detection

[0018] In general, the drowsiness onset detection implementations described herein can predict when a person transitions from a state of wakefulness to a state of drowsiness based on heart rate information over time. Appropriate action can then be taken to stimulate the person to a state of wakefulness or notify other people of their state (with respect to drowsiness/alertness). This is useful in many scenarios where a person's wakefulness is required, such as while driving a car, commandeering a ship, piloting an aircraft, people studying for exams, during meetings, and so on.

[0019] The drowsiness onset detection implementations described herein generally involve capturing a person's heart rate information over time using one or more heart rate (HR) sensors and then computing a heart-rate variability (HRV) signal from the captured heart rate information. This HRV signal is analyzed to extract features that are indicative of an individual's transition from a wakeful state to a drowsy state (i.e., N1 NREM sleep). The extracted features are input into an artificial neural net (ANN) that has been trained using the same features to identify when an individual makes the aforementioned transition to drowsiness. Whenever an onset of drowsiness is detected, a warning is initiated that can take various forms.

[0020] An advantageous feature of the drowsiness onset detection described herein is that the features extracted from the HRV signal have been tailored to specifically detect the transition from wakefulness to drowsiness. For example, in one implementation, just 11 features are employed as opposed to the many features that can be extracted from a HRV signal. As such, the amount of processing, as well as the time it takes to detect the onset of drowsiness, is significantly reduced.

[0021] The drowsiness onset detection implementations described herein are also advantageous for various other reasons including, but not limited to, the following. As will be appreciated from the more-detailed description that follows, some of the drowsiness

onset detection implementations are self-contained. More particularly, they are either fully contained in a local HR sensor device, or employ such a device to capture HR information which is then forwarded to a computing device carried by the user. As such, outside communication is not needed (e.g., no cellular data usage), and the scheme can function even without outside connectivity (for example where cellular service is not available). In addition, the drowsiness onset detection implementations described herein are noninvasive and produce accurate results that are believed to exceed any wakefulness monitoring scheme currently being employed.

1.1 System and Process Frameworks

[0022] This section describes different exemplary implementations of a system framework and a process framework that can be used to realize the drowsiness onset detection implementations described herein. It is noted that in addition to the system framework and process framework implementations described in this section, various other system framework and process framework implementations may also be used to realize the drowsiness onset detection implementations.

[0023] FIG. 1 illustrates one implementation, in simplified form, of a system framework for realizing the drowsiness onset detection implementations described herein. As exemplified in FIG. 1, the system framework 100 includes a set of one or more heart rate (HR) sensors 102. Any type of sensor that captures the heart rate of an individual can be employed as the heart rate (HR) sensor(s) 102. In one version, these HR sensor(s) 102 are physically disposed on (e.g., worn on) the body of a user. In another version, the HR sensor(s) 102 are carried by the user 104. In yet another version, the HR sensor(s) 102 are remote from the user, but still able to detect the person's heart rate (e.g., via imaging techniques). As will be appreciated from the more detailed description that follows, the set of HR sensors 102 is configured to continuously (e.g., on an ongoing basis) and passively measure (e.g., capture) heart rate information associated with the user as they go about their day, and at least output a time-stamped data stream that includes a current value of this rate information.

[0024] Referring again to FIG. 1, the system framework 100 also includes one or more computing devices 104. In one version, a computing device 104 is resident in the HR sensor or sensors 102, thus creating a unified drowsiness onset detection device. In another version, a computing device 104 separate from the HR sensor or sensors 102 is employed. This latter version of the computing device 104 can be a mobile device that is carried by the user, such as a conventional smartphone or a conventional tablet computer. In yet

another version, there is both a computing device 104 resident in the HR sensor(s) 102 and a separate computing device carried by the user. As will be described in more detail later, the computing device 104 resident in the sensor(s) 102 in this last version is used to perform at least some pre-processing of the signal(s) output by the sensor(s). For example,
5 as will be described in more detail later, this pre-processing can involve deriving a heart rate variability (HRV) signal from the raw HR sensor readings.

[0025] As indicated above, the HR sensor or sensors 102 can take different forms. For example, in one implementation where the sensor(s) are worn on the body of a user, the HR sensor(s) 102 can be packaged in wearable device. One such wearable device is a
10 health/fitness tracking device that is worn on an individual's wrist. It is noted that many different types of health/fitness tracking wearable computing devices are commercially available today that include HR sensors.

[0026] In the foregoing versions having a computing device 104 that is separate from the HR sensor or sensors 102, each of the HR sensors is configured to wirelessly transmit the
15 time-stamped data stream output from the sensor(s) to the computing device. The computing device 104 is accordingly configured to wirelessly receive the various data streams transmitted from the set of HR sensors 102. The wireless communication of the various data streams output from the set of HR sensors 102 can be realized using various wireless technologies. For example, this wireless communication can be realized using a
20 conventional Bluetooth personal area network. Another version is possible where the wireless communication is realized using a conventional Wi-Fi local area network. Yet another version is also possible where the wireless communication is realized using a combination of different wireless networking technologies.

[0027] As exemplified in FIG. 1, the system framework 100 also includes a drowsiness
25 onset detector computer program 106 having sub-programs executable by one or more computing devices 104. The sub-programs include a receiving sub-program 108, a feature extraction sub-program 110, a feature combining sub-program 112, an artificial neural network (ANN) classifier input sub-program 114, an onset of drowsiness identifying sub-program 116 and a drowsiness onset warning sub-program 118. In one version, raw HR
30 sensor data 120 is input into the receiving sub-program 108 from the HR sensors 102. In another version, the aforementioned HRV signal 122 is input into the receiving sub-program 108 from the HR sensors 102. The alternate nature of these two inputs is indicated in Fig. 1 by the use of broken line arrows. A drowsiness onset warning 124 is output from the drowsiness onset warning sub-program 118. As indicated previously, each

of these sub-programs is realized on the one or more computing devices (such as those described in more detail in the Exemplary Operating Environments section which follows). It is noted that whenever there is a plurality of computing devices they are in communication with each other via a computer network (such as the Internet or a proprietary intranet).

[0028] Referring now to FIG. 2, the aforementioned one or more computing devices are directed by the foregoing sub-programs of the computer program to accomplish a series of process actions. More particularly, FIG. 2 illustrates an exemplary implementation, in simplified form, of a process for detecting the onset of drowsiness for an individual. This process begins with receiving the HR information from the heart rate sensor or sensors (process action 200). A set of features are then extracted from the HR information (process action 202). These features, which are among many features that can be extracted from HR information, have been determined to be specifically indicative of a transition from a wakeful state of an individual to a drowsy state of an individual. Next, the extracted features are combined to produce a drowsiness detection input (process action 204). The drowsiness detection input is input into an artificial neural network (ANN) classifier that has been trained to distinguish between the wakeful state of an individual and the drowsy state of an individual based on the extracted features (process action 206). It is then identified from an output of the ANN classifier if the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness (process action 208). If so, a drowsiness onset warning is initiated (process action 210). If not, the foregoing process is repeated.

[0029] With regard to initiating the drowsiness onset warning, the warnings can be characterized as belonging to three general categories. The first category is that of a wake-up call to the individual being monitored for drowsiness. By way of example but not limitation, the drowsiness onset warning can be a "message" that is displayed on a display screen. When the heart rate sensor or sensors are packaged such as described previously, this package can have a display. In addition, for implementations employing the aforementioned mobile device, this device can have a display. The message can be text or images, or even take the form of flashing images that appear to the user as a flashing light on the display screen. In addition to, or instead of the foregoing wake-up message examples, the wake-up message can take the form of suggestions to the individual. For example, in the context of an individual who is driving a vehicle, the message could include text providing the location of the closest hotel or rest stop. The warning can also

be an audible alert (e.g., buzzing, ringing, speech, music, and so on) that is played over speaker's resident in the heart rate sensor(s) package or the mobile device. The warning can further be a haptic alert (e.g., vibration, mild electrical shock, and so on) that is created using conventional methods by the heart rate sensor(s) package or the mobile device in response to an initiating command. Any of the foregoing alerts can also be combined to create the drowsiness warning to the individual. Another type of warning involves transmitting a message to a third party (e.g., nurse, medical professional, doctor, sleep researcher, and so on) who in turn warns the individual being monitored for drowsiness, or otherwise takes action based on the individual experiencing the onset of drowsiness. Yet another type of warning is an instruction to a device to take some action in response to the individual experiencing the onset of drowsiness. For example, in the near future the individual may be driving a so-called connected car which can accept an instruction (assuming the proper security protocols are followed) to take over control of the car and drive it to a safe place (like the side of the road). Similarly, other "connected" devices could be controlled in a similar manner to avoid accidents owing to the individual's drowsiness. For example, home appliances like a stove could be turned off.

[0030] As indicated previously, the foregoing implementations and versions of drowsiness onset detection have the advantage of being self-contained. They are either fully contained in a HR sensing device, or employ such a device to capture HR information which is then forwarded to a computing device carried by the user. However, while being self-contained has its advantages, other considerations may outweigh them. For example, in a cloud based scenario, additional historical features such as a person's sleep history, current weather and other realtime info can be included. Since processing power on the cloud is typically quite large, one can conceive a more complex model that involves crunching the sleep info of multiple people at the same time. To this end, some implementations of drowsiness onset detection involve using the foregoing system elements to communicate with a server or cloud service. In this case, the raw HR sensor information or a pre-processed HRV signal is sent to the server or cloud service which would in turn detect when the onset of drowsiness occurs and issue the appropriate warning. This could involve sending a warning indicator back to the user's device(s) which would implement the warning, or to a third party, or both.

[0031] FIG. 3 illustrates the foregoing alternate implementation. As exemplified in FIG. 3, the system framework 300 includes a set of one or more heart rate (HR) sensors 302. As in the previous implementations, in one version, these HR sensor(s) 302 are physically

disposed on (e.g., worn on) the body of a user. In another version, the HR sensor(s) 302 are carried by the user. In yet another version, the HR sensor(s) 302 are remote from the user, but still able to detect the person's heart rate. The set of HR sensors 302 is configured to continuously (e.g., on an ongoing basis) and passively measure (e.g., capture) heart rate information associated with the user as they go about their day, and at least output a time-stamped data stream that includes a current value of this rate information. The system framework 300 also includes one or more computing devices 304. As with previous implementations, in one version, a computing device 304 is resident in the HR sensor or sensors 302. In another version, a computing device 304 separate from the HR sensor or sensors 302 is employed. This latter version of the computing device 304 can be a mobile device that is carried by the user, such as a conventional smartphone or a conventional tablet computer. In yet another version, there is both a computing device 304 resident in the HR sensor(s) 302 and a separate computing device carried by the user. As will be described in more detail later, the computing device 304 resident in the sensor(s) 302 in this last version is used to perform at least some pre-processing of the signal(s) output by the sensor(s), such as deriving the aforementioned HRV from the raw HR sensor readings. Additionally, as before, in the foregoing versions having a computing device 304 that is separate from the HR sensor or sensors 302, each of the HR sensors is configured to wirelessly transmit the time-stamped data stream output from the sensor(s) to the computing device. The computing device 304 is accordingly configured to wirelessly receive the various data streams transmitted from the set of HR sensors 302. The wireless communication of the various data streams output from the set of HR sensors 302 can be realized using various wireless technologies. For example, this wireless communication can be realized using a conventional Bluetooth personal area network. Another version is possible where the wireless communication is realized using a conventional Wi-Fi local area network. Yet another version is also possible where the wireless communication is realized using a combination of different wireless networking technologies.

[0032] Additionally, in the implementations illustrated in Fig. 3, the computing device 304 (whether resident only in the HR sensor(s) 302 or in the form of a separated device carried by the user) is further configured to communicate over a data communication network 306 such as the Internet (among other types of networks) to a cloud service 308 that operates on one or more other computing devices 310 that are remotely located from the computing device 304. The remote computing devices 310 can also communicate with each other via the network 306. The term "cloud service" is used herein to refer to a web

application that operates in the cloud and can be hosted on (e.g., deployed at) a plurality of data centers that can be located in different geographic regions (e.g., different regions of the world). It is noted that in one version of the foregoing scheme the cloud service is replaced by a conventional remote server (not shown).

5 **[0033]** As further exemplified in FIG. 3, the system framework 300 includes a drowsiness onset detector computer program 312 operating on the one or more computing devices 310 associated with the cloud service 308 (or remote server). The drowsiness onset detector computer program 312 has sub-programs executable by the one or more computing devices 310. The sub-programs include a receiving sub-program 314, a feature extraction
10 sub-program 316, a feature combining sub-program 318, an artificial neural network (ANN) classifier input sub-program 320, an onset of drowsiness identifying sub-program 322 and a drowsiness onset warning sub-program 324. In one version, raw HR sensor data 326 is received from the computing device 304 is input into the receiving sub-program 314. In another version, the aforementioned HRV signal 328 is input into the receiving
15 sub-program 314. The alternate nature of these two inputs is indicated in Fig. 3 by the use of broken line arrows. A drowsiness onset warning 330 is output via the network 306 from drowsiness onset warning sub-program 324. As indicated previously, each of these sub-programs is realized on the one or more computing devices (such as those described in more detail in the Exemplary Operating Environments section which follows). It is noted
20 that whenever there is a plurality of computing devices they are in communication with each other via a computer network (such as the Internet or a proprietary intranet).

[0034] Referring now to FIG. 4, the aforementioned one or more computing devices associated with the cloud service (or remote server) are directed by the foregoing sub-programs of the drowsiness onset detector computer program to accomplish a series of
25 process actions. More particularly, FIG. 4 illustrates an exemplary implementation, in simplified form, of a process for detecting the onset of drowsiness for an individual. The process begins with receiving HR information of the individual (process action 400). This HR information is received from the computing device associated with the aforementioned one or more HR sensors via the data communication network. A set of features are then
30 extracted from the HR information (process action 402). These features, which are among many features that can be extracted from HR information, have been determined to be specifically indicative of a transition from a wakeful state of an individual to a drowsy state of an individual. Next, the extracted features are combined to produce a drowsiness detection input (process action 404). The drowsiness detection input is input into an

artificial neural network (ANN) classifier that has been trained to distinguish between the wakeful state of an individual and the drowsy state of an individual based on the extracted features (process action 406). It is then identified from an output of the ANN classifier if the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness (process action 408). If so, a drowsiness onset notification is transmitted (process action 410). If not, the foregoing process is repeated.

[0035] As described previously, the self-contained implementations (such as illustrated in Figs. 1 and 2) initiate a drowsiness onset warning. The drowsiness onset notification transmitted by the cloud service (or remote server) is used for a similar effect. As indicated previously one of type of warning involves a wake-up call to the individual being monitored for drowsiness. Namely, by way of example but not limitation, the drowsiness onset notification can be a "message" that is transmitted to the previously described packaged heart rate sensor or sensors or the aforementioned mobile device. This message can be text or images that are displayed on a display screen associated with the sensor(s) or mobile device, or even take the form of flashing images that appear to the user as a flashing light on the display. In addition to, or instead of the foregoing wake-up message examples, the wake-up message can take the form of suggestions to the individual. For example, in the context of an individual who is driving a vehicle, the message could include text providing the location of the closest hotel or rest stop. The message can also take the form of an instruction to create an audible alert (e.g., buzzing, ringing, speech, music, and so on) that is played over speakers resident in the sensor(s) or mobile device. The message can further take the form of an instruction to create a haptic alert (e.g., vibration, mild electrical shock, and so on) that is created using conventional methods by the sensor(s) or mobile device. As before, any of the foregoing alerts can be combined to create the drowsiness warning to the individual. Another type of notification involves transmitting a message to a third party (e.g., nurse, medical professional, doctor, sleep researcher, and so on) who in turn warns the individual being monitored for drowsiness, or otherwise takes action based on the individual experiencing the onset of drowsiness. Yet another type of notification that the cloud service (or server) could transmit is an instruction to a device to take some action in response to the individual experiencing the onset of drowsiness. For example, in the near future the individual may be driving a so-called connected car which can accept an instruction (assuming the proper security protocols are followed) to take over control of the car and drive it to a safe place (like the side of the road). Similarly, other "connected" devices could be controlled in a

similar manner to avoid accidents owing to the individual's drowsiness. For example, home appliances like a stove could be turned off.

1.2 Feature Extraction

[0036] As indicted previously, in either the self-contained or cloud service-based implementations, an artificial neural network (ANN) classifier that has been trained to distinguish between the wakeful state of an individual and the drowsy state of an individual based on the aforementioned extracted features is employed to detect the onset of drowsiness in an individual being monitored. In this section, exemplary feature extraction methods and exemplary types of features that are extracted will be described.

[0037] Medical research indicates that heart rate and an individual's perception of drowsiness are connected. More particularly, it is possible to detect when an individual transition from a wakeful state to a drowsy state (i.e., N1 NREM sleep) using a measure of that person's heart rate. The wave pattern produced by a typical electrocardiogram (EKG or ECG) includes a series of QRS wave complexes. Each of these complexes is made up of a combination of three deflections—namely the deflections referred to as the Q-wave, R-wave and S-wave. The Q-wave and S-wave are typically characterized by an initial negative or downward deflection, while the R-wave is typically characterized by an initial positive or upward deflection. The time elapsed between two consecutive R waves in the electrocardiogram is referred to as an RR interval. The HR sensor or sensors employed in drowsiness onset detection implementations described herein are capable of at least detecting the R-wave peaks. In one implementation, the detected R-wave peaks forms the aforementioned raw HR signal captured by the HR sensor or sensors. Heart rate variability (HRV) is the physiological phenomenon of variation in the time interval between the R-wave peaks. It is measured by the variation in the RR interval. The HRV signal is computed from the raw HR signal either by a computing device associated with the HR sensor(s), or a separate computing device (such as a mobile device carried by the individual being monitored, or by computing device(s) associated with a cloud service (or remote server)).

[0038] In one implementation, as the HRV signal is generated, it is segmented into a sequence of prescribed-length segments. In one version, these segments are two minutes in length and represent approximately 120 RR interval values (i.e., a 1 Hz signal segment). In one implementation, the two minute segments of the HRV signal are then up-sampled prior to feature extraction. More particularly, curve fitting techniques (such cubic spline interpolation) are employed to model the variation seen in the RR interval values of the

segment as an equation. The resulting equations are then sampled. In one version, where 120 RR intervals are represented by the equations, 840 samples are taken using conventional methods. This represents a seven fold up-sampling of the original HRV segment. This up-sampling (i.e., 1Hz to 7Hz) is performed as it was found a 7 Hz signal
5 segment operates well with the feature extraction methods to be described shortly.

[0039] In one implementation, the aforementioned prescribed-length HRV signal segments are taken from the HRV signal on a rolling window basis such that there is a prescribed offset time between each segment. For example, in the case of a two minute long segment, the offset can be any time period greater than 0 and less than 2 minutes.
10 In one version, a 1 minute offset is employed such that 50% of the data is common with the previous segment and 50 % is new data.

[0040] The aforementioned set of features is extracted from each consecutive segment of the up-sampled HRV signal. In one implementation, this feature extraction generally involves the use of a Discrete Fast Fourier Transform (DFFT), and a Discrete Wavelet Transform (DWT) employing a set of Symlet mother waves of order 3. While many
15 characteristics can be derived from the resulting frequency transformed HRV signal segment, it has been found that the ratio of the power of the low frequency (LF) components to the power of the high frequency (HF) components is particularly indicative of the transition between wakefulness and drowsiness. As such, in one implementation,
20 this LF to HF power ratio is chosen as one of the aforementioned extracted features.

[0041] With regard to the DWT, in one implementation, an 8-level decomposition of the HRV segment is performed. The entropy of a difference between the decomposed segment and the Symlet mother waves at each level is computed using conventional methods to produce a series of upper range D entropy coefficients (i.e., $D_1, \dots, D_8$) and a series of
25 lower range A entropy coefficients (i.e., $A_1, \dots, A_8$). Various averages of the coefficients can also be computed and employed as extracted features, among other combinations. However, it has been found that some of the DWT coefficients and combinations are more indicative of the transition between wakefulness and drowsiness, than others. For example, it has been found that the following DWT coefficients are more indicative of the transition
30 than the others:

- a) D_1 through D_8 entropy coefficients;
- b) A_8 entropy coefficient; and
- c) the mean of the A_1 through A_8 entropy coefficients

[0042] Accordingly, of the many coefficients that can be extracted using DFFT and DWT analyses of the HRV segments, the eleven coefficients described above (i.e., LF/HF power ratio, D₁ through D₈ entropy, A₈ entropy, and A₁ through A₈ entropy mean) were found to provide a specific indication of the transition between wakefulness and drowsiness.

5 1.3 Classifier Training

[0043] FIG. 5 illustrates an exemplary implementation, in simplified form, of a system framework used in training a drowsiness onset detection classifier which in turn can be used to detect the onset of drowsiness in an individual. As exemplified in FIG. 5, the system framework 500 includes a drowsiness onset detector training computer program 502 having sub-programs executable by one or more computing devices. These sub-programs include a receiving sub-program 504, and for the HR information associated with each of a plurality of individuals (I₁, I₂, ..., I_n), an extracting sub-program 506 and a combining sub-program 508. In addition, there is a training sub-program 510 that is trained using the output of the combining sub-program for each individual (I₁, I₂, ..., I_n). HR information 512 output by one or more heart rate sensors for a plurality of individuals, as well as a drowsiness indicator 514 for each of the individuals specifying whether that individual is in a wakeful state or in a drowsy state, are input to the receiving sub-program 504. A trained drowsiness onset detection classifier 516 is output from the training sub-program 510. Each of these sub-programs is realized on one or more computing devices such as that described in more detail in the Exemplary Operating Environments section which follows. It is noted that whenever there is a plurality of computing devices they can be in communication with each other via a computer network (such as the Internet or a proprietary intranet).

[0044] Referring now to FIG. 6, the aforementioned one or more computing devices are directed by the foregoing sub-programs of the drowsiness onset detector training computer program to accomplish a series of process actions. More particularly, HR information output by one or more heart rate sensors for a plurality of individuals, as well as a drowsiness indicator for each of the individuals specifying whether that individual is in a wakeful state or in a drowsy state at the time the HR information was captured, are received via the receiving sub-program (process action 600). In one implementation, the drowsiness indicator is derived from an EEG signal captured using the previously-described electroencephalography methods for identifying whether a person is in a wakeful state of a drowsy state (i.e., N1 NREM sleep). Further, the HR information

captured and the associated drowsiness indicator correspond to times the individual was transitioning from a wakeful state to a drowsy state.

[0045] HR information associated with a previously unselected one of the aforementioned individuals is selected next (process action 602). A set of features is extracted from the selected HR information using the extracting sub-program (process action 604). These features, which are among many features that can be extracted from HR information, have been determined to be specifically indicative of a transition from a wakeful state of an individual to a drowsy state of an individual. The extracted features are then combined to produce a drowsiness detection input using the combining sub-program (process action 606). It is next determined if there is HR information associated with an individual that has not yet been selected and processed (process action 608). If there is, then process actions 602 through 608 are repeated. When the HR information associated with all the individuals has been selected and processed, the training sub-program then trains an artificial neural network (ANN) classifier to distinguish between the wakeful state of an individual and the drowsy state of an individual using the drowsiness detection inputs and drowsiness indicator associated with each of the plurality of individuals. (process action 610). This training is accomplished using conventional ANN training methods.

2.0 Exemplary Operating Environments

[0046] The drowsiness onset detection implementations described herein are operational using numerous types of general purpose or special purpose computing system environments or configurations. FIG. 7 illustrates a simplified example of a general-purpose computer system with which various aspects and elements of drowsiness onset detection, as described herein, may be implemented. It is noted that any boxes that are represented by broken or dashed lines in the simplified computing device 10 shown in FIG. 7 represent alternate implementations of the simplified computing device. As described below, any or all of these alternate implementations may be used in combination with other alternate implementations that are described throughout this document. The simplified computing device 10 is typically found in devices having at least some minimum computational capability such as personal computers (PCs), server computers, handheld computing devices, laptop or mobile computers, communications devices such as cell phones and personal digital assistants (PDAs), multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, and audio or video media players.

[0047] To realize the drowsiness onset detection implementations described herein, the device should have a sufficient computational capability and system memory to enable basic computational operations. In particular, the computational capability of the simplified computing device 10 shown in FIG. 7 is generally illustrated by one or more processing unit(s) 12, and may also include one or more graphics processing units (GPUs) 14, either or both in communication with system memory 16. Note that that the processing unit(s) 12 of the simplified computing device 10 may be specialized microprocessors (such as a digital signal processor (DSP), a very long instruction word (VLIW) processor, a field-programmable gate array (FPGA), or other micro-controller) or can be conventional central processing units (CPUs) having one or more processing cores.

[0048] In addition, the simplified computing device 10 may also include other components, such as, for example, a communications interface 18. The simplified computing device 10 may also include one or more conventional computer input devices 20 (e.g., touchscreens, touch-sensitive surfaces, pointing devices, keyboards, audio input devices, voice or speech-based input and control devices, video input devices, haptic input devices, devices for receiving wired or wireless data transmissions, and the like) or any combination of such devices.

[0049] Similarly, various interactions with the simplified computing device 10 and with any other component or feature of wearable sensing, including input, output, control, feedback, and response to one or more users or other devices or systems associated with drowsiness onset detection, are enabled by a variety of Natural User Interface (NUI) scenarios. The NUI techniques and scenarios include, but are not limited to, interface technologies that allow one or more users user to interact in a “natural” manner, free from artificial constraints imposed by input devices such as mice, keyboards, remote controls, and the like.

[0050] Such NUI implementations are enabled by the use of various techniques including, but not limited to, using NUI information derived from user speech or vocalizations captured via microphones or other sensors. Such NUI implementations are also enabled by the use of various techniques including, but not limited to, information derived from a user’s facial expressions and from the positions, motions, or orientations of a user’s hands, fingers, wrists, arms, legs, body, head, eyes, and the like, where such information may be captured using various types of 2D or depth imaging devices such as stereoscopic or time-of-flight camera systems, infrared camera systems, RGB (red, green and blue) camera systems, and the like, or any combination of such devices. Further examples of such NUI

implementations include, but are not limited to, NUI information derived from touch and stylus recognition, gesture recognition (both onscreen and adjacent to the screen or display surface), air or contact-based gestures, user touch (on various surfaces, objects or other users), hover-based inputs or actions, and the like. Such NUI implementations may also include, but are not limited, the use of various predictive machine intelligence processes that evaluate current or past user behaviors, inputs, actions, etc., either alone or in combination with other NUI information, to predict information such as user intentions, desires, and/or goals. Regardless of the type or source of the NUI-based information, such information may then be used to initiate, terminate, or otherwise control or interact with one or more inputs, outputs, actions, or functional features of the drowsiness onset detection implementations described herein.

[0051] However, it should be understood that the aforementioned exemplary NUI scenarios may be further augmented by combining the use of artificial constraints or additional signals with any combination of NUI inputs. Such artificial constraints or additional signals may be imposed or generated by input devices such as mice, keyboards, and remote controls, or by a variety of remote or user worn devices such as accelerometers, electromyography (EMG) sensors for receiving myoelectric signals representative of electrical signals generated by user's muscles, heart-rate monitors, galvanic skin conduction sensors for measuring user perspiration, wearable or remote biosensors for measuring or otherwise sensing user brain activity or electric fields, wearable or remote biosensors for measuring user body temperature changes or differentials, and the like. Any such information derived from these types of artificial constraints or additional signals may be combined with any one or more NUI inputs to initiate, terminate, or otherwise control or interact with one or more inputs, outputs, actions, or functional features of the drowsiness onset detection implementations described herein.

[0052] The simplified computing device 10 may also include other optional components such as one or more conventional computer output devices 22 (e.g., display device(s) 24, audio output devices, video output devices, devices for transmitting wired or wireless data transmissions, and the like). Note that typical communications interfaces 18, input devices 20, output devices 22, and storage devices 26 for general-purpose computers are well known to those skilled in the art, and will not be described in detail herein.

[0053] The simplified computing device 10 shown in FIG. 7 may also include a variety of computer-readable media. Computer-readable media can be any available media that can

be accessed by the computer 10 via storage devices 26, and can include both volatile and nonvolatile media that is either removable 28 and/or non-removable 30, for storage of information such as computer-readable or computer-executable instructions, data structures, programs, sub-programs, or other data. Computer-readable media includes
5 computer storage media and communication media. Computer storage media refers to tangible computer-readable or machine-readable media or storage devices such as digital versatile disks (DVDs), blu-ray discs (BD), compact discs (CDs), floppy disks, tape drives, hard drives, optical drives, solid state memory devices, random access memory (RAM), read-only memory (ROM), electrically erasable programmable read-only memory
10 (EEPROM), CD-ROM or other optical disk storage, smart cards, flash memory (e.g., card, stick, and key drive), magnetic cassettes, magnetic tapes, magnetic disk storage, magnetic strips, or other magnetic storage devices. Further, a propagated signal is not included within the scope of computer-readable storage media.

[0054] Retention of information such as computer-readable or computer-executable
15 instructions, data structures, programs, sub-program, and the like, can also be accomplished by using any of a variety of the aforementioned communication media (as opposed to computer storage media) to encode one or more modulated data signals or carrier waves, or other transport mechanisms or communications protocols, and can include any wired or wireless information delivery mechanism. Note that the terms
20 “modulated data signal” or “carrier wave” generally refer to a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. For example, communication media can include wired media such as a wired network or direct-wired connection carrying one or more modulated data signals, and wireless media such as acoustic, radio frequency (RF), infrared, laser, and other wireless
25 media for transmitting and/or receiving one or more modulated data signals or carrier waves.

[0055] Furthermore, software, programs, and/or computer program products embodying some or all of the various drowsiness onset detection implementations described herein, or portions thereof, may be stored, received, transmitted, or read from any desired
30 combination of computer-readable or machine-readable media or storage devices and communication media in the form of computer-executable instructions or other data structures. Additionally, the claimed subject matter may be implemented as a method, apparatus, or article of manufacture using standard programming and/or engineering techniques to produce software, firmware, hardware, or any combination thereof to control

a computer to implement the disclosed subject matter. The term “article of manufacture” as used herein is intended to encompass a computer program accessible from any computer-readable device, or media.

[0056] The drowsiness onset detection implementations described herein may be further described in the general context of computer-executable instructions, such as programs, being executed by a computing device. Generally, programs include routines, programs, objects, components, data structures, and the like, that perform particular tasks or implement particular abstract data types. The drowsiness onset detection implementations described herein may also be practiced in distributed computing environments where tasks are performed by one or more remote processing devices, or within a cloud of one or more devices, that are linked through one or more communications networks. In a distributed computing environment, programs may be located in both local and remote computer storage media including media storage devices. Additionally, the aforementioned instructions may be implemented, in part or in whole, as hardware logic circuits, which may or may not include a processor.

[0057] Alternatively, or in addition, the functionality described herein can be performed, at least in part, by one or more hardware logic components. For example, and without limitation, illustrative types of hardware logic components that can be used include field-programmable gate arrays (FPGAs), application-specific integrated circuits (ASICs), application-specific standard products (ASSPs), system-on-a-chip systems (SOCs), complex programmable logic devices (CPLDs), and so on.

3.0 Other Implementations

[0058] It is noted that any or all of the aforementioned implementations throughout the description may be used in any combination desired to form additional hybrid implementations. In addition, although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described above. Rather, the specific features and acts described above are disclosed as example forms of implementing the claims.

[0059] What has been described above includes example implementations. It is, of course, not possible to describe every conceivable combination of components or methodologies for purposes of describing the claimed subject matter, but one of ordinary skill in the art may recognize that many further combinations and permutations are possible.

Accordingly, the claimed subject matter is intended to embrace all such alterations, modifications, and variations that fall within the spirit and scope of the appended claims.

[0060] In regard to the various functions performed by the above described components, devices, circuits, systems and the like, the terms (including a reference to a "means") used to describe such components are intended to correspond, unless otherwise indicated, to any component which performs the specified function of the described component (e.g., a functional equivalent), even though not structurally equivalent to the disclosed structure, which performs the function in the herein illustrated exemplary aspects of the claimed subject matter. In this regard, it will also be recognized that the foregoing implementations include a system as well as a computer-readable storage media having computer-executable instructions for performing the acts and/or events of the various methods of the claimed subject matter.

[0061] There are multiple ways of realizing the foregoing implementations (such as an appropriate application programming interface (API), tool kit, driver code, operating system, control, standalone or downloadable software object, or the like), which enable applications and services to use the implementations described herein. The claimed subject matter contemplates this use from the standpoint of an API (or other software object), as well as from the standpoint of a software or hardware object that operates according to the implementations set forth herein. Thus, various implementations described herein may have aspects that are wholly in hardware, or partly in hardware and partly in software, or wholly in software.

[0062] The aforementioned systems have been described with respect to interaction between several components. It will be appreciated that such systems and components can include those components or specified sub-components, some of the specified components or sub-components, and/or additional components, and according to various permutations and combinations of the foregoing. Sub-components can also be implemented as components communicatively coupled to other components rather than included within parent components (e.g., hierarchical components).

[0063] Additionally, it is noted that one or more components may be combined into a single component providing aggregate functionality or divided into several separate sub-components, and any one or more middle layers, such as a management layer, may be provided to communicatively couple to such sub-components in order to provide integrated functionality. Any components described herein may also interact with one or

more other components not specifically described herein but generally known by those of skill in the art.

4.0 Claim Support And Further Implementations

[0064] The following paragraphs summarize various examples of implementations which may be claimed in the present document. However, it should be understood that the implementations summarized below are not intended to limit the subject matter which may be claimed in view of the foregoing descriptions. Further, any or all of the implementations summarized below may be claimed in any desired combination with some or all of the implementations described throughout the foregoing description and any implementations illustrated in one or more of the figures, and any other implementations described below. In addition, it should be noted that the following implementations are intended to be understood in view of the foregoing description and figures described throughout this document.

[0065] In one implementation, a system is employed for detecting the onset of drowsiness in an individual. This system includes one or more heart rate (HR) sensors which capture HR information of the individual, and one or more computing devices, the computing devices being in communication with each other via a computer network whenever there is a plurality of computing devices. The system also includes a drowsiness onset detector computer program having a plurality of sub-programs executable by the one or more computing devices. The one or more computing devices are directed by the sub-programs of the drowsiness onset detector computer program to receive the HR information from the heart rate sensor or sensors, extract a set of features from the HR information (which are among many features that can be extracted from HR information) that have been determined to be specifically indicative of a transition from a wakeful state of an individual to a drowsy state of an individual, combine the extracted features to produce a drowsiness detection input, input the drowsiness detection input into an artificial neural network (ANN) classifier that has been trained to distinguish between the wakeful state of an individual and the drowsy state of an individual based on the extracted features, identify from an output of the ANN classifier if the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, and whenever the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, initiate a drowsiness onset warning.

[0066] In one implementation of the system where the HR information includes a raw HR signal captured by the one or more HR sensors, the sub-program for extracting the set

of features from the HR information includes first computing a heart rate variability (HRV) signal from the raw HR signal. In another implementation of the system, the one or more HR sensors includes a computing device which computes a heart rate variability (HRV) signal from a raw HR signal captured by the HR sensor or sensors. The

5 aforementioned HR information includes this HRV signal. In yet another implementation, the one or more heart rate (HR) sensors and the one or more computing devices are resident in a unified drowsiness onset detection device. In one version of this implementation, the drowsiness onset detection device is either a wearable device worn on the person of the individual, or a mobile device carried by the individual. In still another

10 implementation, the one or more heart rate (HR) sensors are resident in a wearable device worn on the person of the individual, and the one or more computing devices are resident in a mobile device carried by the individual, and the wearable device is in wireless communication with the mobile device. In one version of this implementation, the HR information includes a raw HR signal captured by the one or more HR sensors which is

15 transmitted to the mobile device carried by the individual. In another version, the HR information includes a raw HR signal captured by the one or more HR sensors, and the one or more HR sensors includes a computing device which computes a heart rate variability (HRV) signal from the raw HR signal, where the HRV signal is transmitted to the mobile device carried by the individual.

20 **[0067]** One implementation of the system further includes a segmenting sub-program that is executed before the execution of the sub-program for extracting the set of features from the HR information. This segmenting sub-program segments the HR information received from the heart rate sensor or sensors into a sequence of prescribed-length segments whenever the received HR information is in the form of a heart rate variability

25 (HRV) signal. In addition, the segmenting sub-program computes a HRV signal from the received HR information whenever the received HR information is not in the form of a heart rate variability (HRV) signal before segmenting the HRV signal into the sequence of prescribed-length segments. The sub-programs for extracting a set of features from the HR information, combining the extracted features, inputting the drowsiness detection input

30 into the ANN classifier, identifying from an output of the ANN classifier if the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, and whenever the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, initiating a drowsiness onset warning, are executed on each of the HRV signal segments as they are created. In one version, these prescribed-length segments of

the HRV signal are each two minutes in length and represent a rolling window such that there is a prescribed offset time between each segment. In one variant, each two minute long HRV signal segment represents approximately 120 heart rate variation values, and the segmenting sub-program further includes up-sampling each HRV signal segment to approximately 840 samples heart rate variation values using curve fitting techniques. In another version, the sub-program for extracting the set of features from each HRV signal segment includes extracting the set of features using a Discrete Fast Fourier Transform (DFFT) and a Discrete Wavelet Transform (DWT). In one variant, the DFFT is used to compute a ratio of the power of the low frequency components of the HRV signal segment to the power of the high frequency (HF) components of the HRV signal segment as an extracted feature, and the DWT is used to compute from an 8-level decomposition of the HRV segment the D_1 through D_8 entropy coefficients, A_8 entropy coefficient and the mean of the A_1 through A_8 entropy coefficients as extracted features.

[0068] In one implementation of the system, the sub-program for initiating a drowsiness onset warning includes initiating a wake-up call to the individual in the form of a message that is displayed on a display screen that is viewable by the individual. In one version, this message includes text, or images, or flashing images that appear as a flashing light on the display screen, or any combination of the foregoing. In another implementation, the sub-program for initiating a drowsiness onset warning includes initiating a wake-up call to the individual in the form of an audible alert that is played over speakers which are within a hearing range of the individual. In yet another implementation, the sub-program for initiating a drowsiness onset warning includes initiating a wake-up call to the individual in the form of a haptic alert that can be felt by the individual. In still another implementation, the sub-program for initiating a drowsiness onset warning includes transmitting a message to a third party that requests the third party to warn the individual.

[0069] In one implementation, another system is employed for detecting the onset of drowsiness in an individual. This system includes one or more computing devices, the computing devices being in communication with each other via a computer network whenever there is a plurality of computing devices. The system also includes a drowsiness onset detector computer program having a plurality of sub-programs executable by the one or more computing devices. The one or more computing devices are directed by the sub-programs of the drowsiness onset detector computer program to receive heart rate (HR) information of the individual, where the one or more computing devices are in communication over a data communication network with a remote computing device

associated with one or more HR sensors which capture the HR information, and where the HR information is received from the remote computing device via the data communication network. The one or more computing devices are also directed by the sub-programs of the drowsiness onset detector computer program to extract a set of features from the HR information (which are among many features that can be extracted from HR information) that have been determined to be specifically indicative of a transition from a wakeful state of an individual to a drowsy state of an individual, combine the extracted features to produce a drowsiness detection input, input the drowsiness detection input into an artificial neural network (ANN) classifier that has been trained to distinguish between the wakeful state of an individual and the drowsy state of an individual based on the extracted features, identify from an output of the ANN classifier if the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, and whenever the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, transmit a drowsiness onset notification.

[0070] In one implementation, another system is employed for training a drowsiness onset detection classifier. This system includes one or more computing devices, the computing devices being in communication with each other via a computer network whenever there is a plurality of computing devices. The system also includes a drowsiness onset detector training computer program having a plurality of sub-programs executable by the one or more computing devices. The one or more computing devices are directed by the sub-programs of the drowsiness onset detector training computer program to receive the HR information output by one or more heart rate sensors for a plurality of individuals, as well as a drowsiness indicator for each of the individuals specifying whether that individual is in a wakeful state or in a drowsy state at the time the HR information was captured. For the HR information associated with each of the plurality of individuals, a set of features is extracted from the HR information (which are among many features that can be extracted from HR information) that have been determined to be specifically indicative of a transition from a wakeful state of an individual to a drowsy state of an individual, and the extracted features are combined to produce a drowsiness detection input. An artificial neural network (ANN) classifier is trained to distinguish between the wakeful state of an individual and the drowsy state of an individual using the drowsiness detection inputs and drowsiness indicator associated with each of the plurality of individuals.

[0071] In one implementation, a computer-implemented process is employed for detecting the onset of drowsiness in an individual, which includes using a computing device to perform the following process actions: receiving the HR information from the heart rate sensor or sensors, extracting a set of features from the HR information (which are among many features that can be extracted from HR information) that have been determined to be specifically indicative of a transition from a wakeful state of an individual to a drowsy state of an individual, combining the extracted features to produce a drowsiness detection input, inputting the drowsiness detection input into an artificial neural network (ANN) classifier that has been trained to distinguish between the wakeful state of an individual and the drowsy state of an individual based on the extracted features, identifying from an output of the ANN classifier if the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, and whenever the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, initiating a drowsiness onset warning.

[0072] In another implementation, a computer-implemented process is employed for detecting the onset of drowsiness in an individual, which includes using a computing device to perform the following process actions: receiving heart rate (HR) information of the individual, where one or more computing devices are in communication over a data communication network with a remote computing device associated with one or more HR sensors which capture the HR information, and where the HR information is received from the remote computing device via the data communication network. Process actions are also included for extracting a set of features from the HR information (which are among many features that can be extracted from HR information) that have been determined to be specifically indicative of a transition from a wakeful state of an individual to a drowsy state of an individual, combining the extracted features to produce a drowsiness detection input, inputting the drowsiness detection input into an artificial neural network (ANN) classifier that has been trained to distinguish between the wakeful state of an individual and the drowsy state of an individual based on the extracted features, identifying from an output of the ANN classifier if the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, and whenever the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, transmitting a drowsiness onset notification.

[0073] In another implementation, a computer-implemented process is employed for training a drowsiness onset detection classifier, which includes using a computing device

to perform the following process actions: receiving the HR information output by one or more heart rate sensors for a plurality of individuals, as well as a drowsiness indicator for each of the individuals specifying whether that individual is in a wakeful state or in a drowsy state at the time the HR information was captured. For the HR information associated with each of the plurality of individuals, extracting a set of features from the HR information (which are among many features that can be extracted from HR information) that have been determined to be specifically indicative of a transition from a wakeful state of an individual to a drowsy state of an individual, and combining the extracted features to produce a drowsiness detection input. An artificial neural network (ANN) classifier is then trained to distinguish between the wakeful state of an individual and the drowsy state of an individual using the drowsiness detection inputs and drowsiness indicator associated with each of the plurality of individuals.

[0074] In one implementation, detecting the onset of drowsiness in an individual includes using a computing device to perform the following process steps: a receiving step for receiving the HR information from the heart rate sensor or sensors, an extracting step for extracting a set of features from the HR information (which are among many features that can be extracted from HR information) that have been determined to be specifically indicative of a transition from a wakeful state of an individual to a drowsy state of an individual, a combining step for combining the extracted features to produce a drowsiness detection input, an inputting step for inputting the drowsiness detection input into an artificial neural network (ANN) classifier that has been trained to distinguish between the wakeful state of an individual and the drowsy state of an individual based on the extracted features, an identifying step for identifying from an output of the ANN classifier if the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, and whenever the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, performing an initiating step for initiating a drowsiness onset warning.

[0075] In another one implementation, detecting the onset of drowsiness in an individual includes using a computing device to perform the following process steps: a receiving step for receiving heart rate (HR) information of the individual, where one or more computing devices are in communication over a data communication network with a remote computing device associated with one or more HR sensors which capture the HR information, and where the HR information is received from the remote computing device via the data communication network. The process steps further include an extracting step

for extracting a set of features from the HR information (which are among many features that can be extracted from HR information) that have been determined to be specifically indicative of a transition from a wakeful state of an individual to a drowsy state of an individual, a combining step for combining the extracted features to produce a drowsiness detection input, an inputting step for inputting the drowsiness detection input into an artificial neural network (ANN) classifier that has been trained to distinguish between the wakeful state of an individual and the drowsy state of an individual based on the extracted features, an identifying step for identifying from an output of the ANN classifier if the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, and whenever the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, performing a transmitting step for transmitting a drowsiness onset notification.

[0076] In another one implementation, training a drowsiness onset detection classifier includes using a computing device to perform the following process steps: a receiving step for receiving the HR information output by one or more heart rate sensors for a plurality of individuals, as well as a drowsiness indicator for each of the individuals specifying whether that individual is in a wakeful state or in a drowsy state at the time the HR information was captured. For the HR information associated with each of the plurality of individuals, performing an extracting step for extracting a set of features from the HR information (which are among many features that can be extracted from HR information) that have been determined to be specifically indicative of a transition from a wakeful state of an individual to a drowsy state of an individual, and a combining step for combining the extracted features to produce a drowsiness detection input. A training step is then performed for training an artificial neural network (ANN) classifier to distinguish between the wakeful state of an individual and the drowsy state of an individual using the drowsiness detection inputs and drowsiness indicator associated with each of the plurality of individuals.

Wherefore, what is claimed is :

1. A system for detecting the onset of drowsiness in an individual, comprising:
 - one or more heart rate (HR) sensors which capture HR information of the individual over time, the HR information comprising heart rate variability (HRV) information;
 - one or more computing devices, said computing devices being in communication with each other whenever there is a plurality of computing devices, and a drowsiness onset detector computer program having a plurality of sub-programs executable by the one or more computing devices, the one or more computing devices being directed by the sub-programs of the drowsiness onset detector computer program to,
 - receive the HR information from the heart rate sensor or sensors,
 - extract a plurality of different HRV-related features from the HR information, said extracted different HRV-related features, which are among many features that can be extracted from HR information, having been determined to be specifically indicative of a transition from a wakeful state of an individual to a drowsy state of an individual,
 - combine the extracted different HRV-related features to produce a drowsiness detection input,
 - input the drowsiness detection input into an artificial neural network (ANN) classifier that has been trained to distinguish between said wakeful state of an individual and said drowsy state of an individual based on the extracted different HRV-related features,
 - identify from an output of the ANN classifier if the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, and
 - whenever the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, initiate a drowsiness onset warning.
- 2 . The system of claim 1, wherein the HR information comprises a raw HR signal captured by the one or more HR sensors, and wherein the sub-program for extracting a plurality of different HRV-related features from the HR information comprises first computing a HRV signal from the raw HR signal.
- 3 . The system of claim 1, wherein the one or more HR sensors comprise a computing device which comprises a HRV signal from a raw HR signal captured by the HR sensor or sensors, and wherein the HR information comprises the HRV signal.
4. The system of claim 1, wherein the one or more heart rate (HR) sensors and the one or more computing devices are resident in a unified drowsiness onset detection device.
5. The system of claim 4, wherein the drowsiness onset detection device is either a wearable device worn on the person of said individual, or a mobile device carried by said individual.

6. The system of claim 1, wherein the one or more heart rate (HR) sensors are resident in a wearable device worn on the person of said individual, and the one or more computing devices are resident in a mobile device carried by said individual, and wherein the wearable device is in wireless communication with the mobile device.

7. The system of claim 6, wherein the HR information comprises a raw HR signal captured by the one or more HR sensors which is transmitted to said mobile device carried by the individual.

8. The system of claim 6, wherein the HR information comprises a raw HR signal captured by the one or more HR sensors, and wherein the one or more HR sensors comprise a computing device which computes a HRV signal from the raw HR signal, said HRV signal being transmitted to said mobile device carried by the individual.

9. The system of Claim 1, further comprising a segmenting sub-program that is executed before the execution of the sub-program for extracting a plurality of different HRV-related features from the HR information, said segmenting sub-program segmenting the HR information received from the heart rate sensor or sensors into a sequence of prescribed-length segments whenever the received HR information is in the form of a HRV signal, and said segmenting sub-program computing a HRV signal from the received HR information whenever the received HR information is not in the form of a HRV signal before segmenting the HRV signal into the sequence of prescribed-length segments, and wherein the sub-programs for extracting a plurality of different HRV-related features from the HR information, combining the extracted different HRV-related features, inputting the drowsiness detection input into the ANN classifier, identifying from an output of the ANN classifier if the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, and whenever the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, initiating a drowsiness onset warning, are executed on each of the HRV signal segments as they are created.

10. The system of claim 9, wherein the prescribed-length segments of the HRV signal are each two minutes in length and represent a rolling window such that there is a prescribed offset time between each segment.

11. The system of claim 10, wherein each two minute long HRV signal segment represents approximately 120 heart rate variation values, and wherein the segmenting subprogram further comprises up-sampling each HRV signal segment to approximately 840 samples heart rate variation values using curve fitting techniques.

12. The system of claim 9, wherein the sub-program for extracting a plurality of different HRV-related features from each HRV signal segment comprises extracting the plurality of different HRV-related features using a Discrete Fast Fourier Transform (DFFT) and a Discrete Wavelet Transform (DWT).

13. The system of claim 12, wherein the DFFT is used to compute a ratio of the power of the low frequency components of the HRV signal segment to the power of the high frequency (HF) components of the HRV signal segment as one of the extracted different HRV-related features, and wherein the DWT is used to compute from an 8-level decomposition of the HRV segment the D_i through D_8 entropy coefficients, A_8 entropy coefficient and the mean of the A_i through A_8 entropy coefficients as others of the extracted different HRV-related features.

14. The system of claim 1, wherein the sub-program for initiating a drowsiness onset warning comprises initiating a wake-up call to the individual in the form of a message that is displayed on a display screen that is viewable by the individual.

15. The system of claim 14, wherein the message comprises at least one of text, or images, or flashing images that appear as a flashing light on the display screen.

16. The system of claim 1, wherein the sub-program for initiating a drowsiness onset warning, comprises initiating a wake-up call to the individual in the form of an audible alert that is played over speakers which are within a hearing range of the individual.

17. The system of claim 1, wherein the sub-program for initiating a drowsiness onset warning, comprises initiating a wake-up call to the individual in the form of a haptic alert that can be felt by the individual.

18. The system of claim 1, wherein the sub-program for initiating a drowsiness onset warning, comprises transmitting a message to a third party requesting the third party to warn the individual.

19. A system for detecting the onset of drowsiness in an individual, comprising:

- one or more computing devices, said computing devices being in communication with each other, whenever there is a plurality of computing devices; and
- a drowsiness onset detector computer program having a plurality of subprograms executable by one or more computing devices, the one or more computing devices being directed by the sub-programs of the drowsiness onset detector computer program to,
 - receive heart rate (HR) information of the individual over time, the HR information comprising heart rate variability (HRV) information, wherein the one or more computing devices are in communication over a data communication network with a remote computing device associated with one or more HR sensors which capture the HR information, and wherein the HR information is received from the remote computing device via the data communication network;
 - extract a plurality of different HRV-related features from the HR information, said extracted different HRV-related features, which are among many features that can be extracted from HR information, having been determined to be specifically indicative of a transition from a wakeful state of an individual to a drowsy state of an individual,
 - combine the extracted different HRV-related features to produce a drowsiness detection input,
 - input the drowsiness detection input into an artificial neural network (ANN) classifier that has been trained to distinguish between said wakeful state of an individual and said drowsy state of an individual based on the extracted different HRV-related features,

identify from an output of the ANN classifier if the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, and

whenever the drowsiness detection input indicates that the individual is exhibiting an onset of drowsiness, transmit a drowsiness onset notification.

20. A drowsiness onset detection classifier training system, comprising :

one or more computing devices, said computing devices being in communication with each other via a computer network whenever there is a plurality of computing devices; and

a drowsiness onset detector training computer program having a plurality of sub-programs executable by the one or more computing devices, the one or more computing devices being directed by the sub-programs of the drowsiness onset detector training computer program to,

receive the HR information output by one or more heart rate sensors for a plurality of individuals over time, as well as a drowsiness indicator for each of the individuals specifying whether that individual is in a wakeful state or in a drowsy state at the time the HR information was captured, the HR information comprising heart rate variability (HRV) information, for the HR information associated with each of the plurality of individuals,

extract a plurality of different HRV-related features from the HR information, said extracted different HRV-related features, which are among many features that can be extracted from HR information, having been determined to be specifically indicative of a transition from a wakeful state of an individual to a drowsy state of an individual, and

combine the extracted different HRV-related features to produce a drowsiness detection input, and train an artificial neural network (ANN) classifier to distinguish between said wakeful state of an individual and said drowsy state of an individual using said the drowsiness detection inputs and drowsiness indicator associated with each of the plurality of individuals.