



(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:  
**23.08.2017 Bulletin 2017/34**

(51) Int Cl.:  
**A61B 5/00 (2006.01)** **A61B 5/103 (2006.01)**  
**A61B 5/024 (2006.01)**

(21) Application number: **17156334.9**

(22) Date of filing: **15.02.2017**

(84) Designated Contracting States:  
**AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR**  
 Designated Extension States:  
**BA ME**  
 Designated Validation States:  
**MA MD**

(72) Inventors:  
 • **Addison, Paul Stanley**  
**Edinburgh, EH10 7BE (GB)**  
 • **Foo, David**  
**Glasgow, G31 2JU (GB)**  
 • **Jacquel, Dominique**  
**Edinburgh, System and Meth (GB)**  
  
 (74) Representative: **Hargreaves, Timothy Edward**  
**Marks & Clerk LLP**  
**Atholl Exchange**  
**6 Canning Street**  
**Edinburgh EH3 8EG (GB)**


(30) Priority: **19.02.2016 US 201662297682 P**  
**13.05.2016 US 201662335862 P**  
**26.09.2016 US 201662399741 P**

(71) Applicant: **Covidien LP**  
**Mansfield, MA 02048 (US)**

(54) **METHODS FOR VIDEO-BASED MONITORING OF VITAL SIGNS**

(57) The present invention relates to the field of medical monitoring, and in particular non-contact, video-based monitoring of pulse rate, respiration rate, motion, and oxygen saturation. Systems and methods are

described for capturing images of a patient, producing intensity signals from the images, filtering those signals to focus on a physiologic component, and measuring a vital sign from the filtered signals.



**FIG.2A**

**Description**

## RELATED APPLICATIONS

**[0001]** This application claims priority to and the benefit of U.S. Provisional Application No. 62/297,682 (filed February 19, 2016); U.S. Provisional Application No. 62/335,862 (filed May 13, 2016); and U.S. Provisional Application No. 62/399,741 (filed September 26, 2016), the contents of which are incorporated herein by reference in their entirety.

## BACKGROUND

**[0002]** Many conventional medical monitors require attachment of a sensor to a patient in order to detect physiologic signals from the patient and transmit detected signals through a cable to the monitor. These monitors process the received signals and determine vital signs such as the patient's pulse rate, respiration rate, and arterial oxygen saturation. An example of a prior art monitoring system 100 is shown in Figure 1. The system 100 includes a monitor 110 and a sensor 112 connected to the monitor 110 by a cable 114. In the example of Figure 1, the monitor 110 is a pulse oximeter, and the sensor 112 is a finger sensor including two light emitters and a photodetector. The sensor 112 emits light into the patient's finger, detects light transmitted through the patient's finger, and transmits the detected light signal through the cable 114 to the monitor 110. The monitor 110 includes a processor that processes the signal, determines vital signs (including pulse rate, respiration rate, and arterial oxygen saturation), and displays them on an integrated display 116.

**[0003]** Other monitoring systems include other types of monitors and sensors, such as electroencephalogram (EEG) sensors, blood pressure cuffs, temperature probes, and others.

**[0004]** Many of these conventional monitors require some type of cable or wire, such as cable 114 in Figure 1, physically connecting the patient to the monitor. As a result, the patient is effectively tethered to the monitor, which can limit the patient's movement around a hospital room, restrict even simple activities such as writing or eating, and prevent easy transfer of the patient to different locations in the hospital without either disconnecting and connecting new monitors, or moving the monitor with the patient.

**[0005]** Some wireless, wearable sensors have been developed, such as wireless EEG patches and wireless pulse oximetry sensors. Although these sensors improve patient mobility, they introduce new problems such as battery consumption, infection risk from re-use on sequential patients, high cost, and bulky designs that detract from patient compliance and comfort.

**[0006]** Video-based monitoring is a new field of patient monitoring that uses a remote video camera to detect physical attributes of the patient. This type of monitoring

may also be called "non-contact" monitoring in reference to the remote video sensor, which does not contact the patient. The remainder of this disclosure offers solutions and improvements in this new field.

## BRIEF DESCRIPTION OF THE DRAWINGS

**[0007]**

Figure 1 is a perspective view of a pulse oximetry monitor and sensor according to the prior art.

Figure 2A is schematic view of a video-based patient monitoring system according to an embodiment of the invention.

Figure 2B is schematic view of a video-based patient monitoring system monitoring multiple patients according to an embodiment of the invention.

Figure 3A depicts an image frame from a video signal according to an embodiment of the invention.

Figure 3B depicts an image frame from a video signal according to an embodiment of the invention.

Figure 4A depicts light intensity signals from the video signal of Figure 3A.

Figure 4B depicts light intensity signals from the video signal of Figure 4A.

Figure 5A depicts an image frame according to an embodiment of the invention.

Figure 5B depicts an image frame according to an embodiment of the invention.

Figure 5C is a chart of a light intensity signal from a first region of interest according to an embodiment of the invention.

Figure 5D is a chart of a light intensity signal from a second region of interest according to an embodiment of the invention.

Figure 5E is a flowchart of a method for measuring a vital sign from a combined region of interest according to an embodiment of the invention.

Figure 5F is a flowchart of a method for dynamically updating and displaying a color signal from a moving region of interest according to an embodiment of the invention.

Figure 6A is a flowchart of a method of determining vital signs from a video signal according to an embodiment of the invention.

Figure 6B is a chart of contact-oximeter-based and video-based vital signs (heart rate and SpO<sub>2</sub>) over time according to an embodiment of the invention.

Figure 7 is a flowchart of a method of calibrating video-based pulse oximetry according to an embodiment of the invention.

Figure 8 is a chart of video-based and contact-based measurements of arterial oxygen saturation over time, including a desaturation event, according to an embodiment of the invention.

Figure 9 is a flowchart of a method for calibrating a video camera according to an embodiment of the invention.

Figure 10A is a chart of red, green, and blue pixel signals over time and a corresponding frequency transform according to an embodiment of the invention.

Figure 10B is a flowchart of a method of calculating heart rate from a video signal utilizing a frequency accumulator, according to an embodiment of the invention.

Figure 11 depicts a patient in an image frame according to an embodiment of the invention.

Figure 12 is a bottom view of a calibration strip according to an embodiment of the invention.

Figure 13 is a set of charts showing three source signals and three independent component signals according to an embodiment of the invention.

Figure 14 is a chart of contact-oximeter-based and video-based ICA-derived heart rate values over time according to an embodiment of the invention.

Figure 15 is a schematic chart illustrating an independent component signal and two source signals, according to an embodiment of the invention.

Figure 16 is a flowchart illustrating a method for utilizing an independent component signal, according to an embodiment of the invention.

Figure 17A is a flowchart illustrating a method for identifying non-physiologic frequencies according to an embodiment of the invention.

Figure 17B is a flowchart illustrating a method for identifying non-physiologic frequencies according to an embodiment of the invention.

Figure 18 is a schematic cut-away view of an optical splitter according to an embodiment of the invention.

Figure 19 is a scatter plot of video-calculated heart rate measurements against reference heart rate measurements, according to an embodiment of the invention.

Figure 20 is a scatter plot of video-based respiration rate measurements against reference respiratory rate measurements, according to an embodiment of the invention.

Figure 21 is a scatter plot of video-based SpO<sub>2</sub> measurements against reference SpO<sub>2</sub> measurements, according to an embodiment of the invention.

## SUMMARY

**[0008]** A calibration strip may be used as explained in more detail below. In an embodiment, a video-based method of measuring a patient's vital sign includes providing a calibration strip comprising a substrate with first and second opposite surfaces, an adhesive on the first surface of the substrate for adhering to a patient, and a visible scale on the second surface for viewing by a video camera; detecting, by the video camera, a first light signal from the scale and a second light signal from the patient, within the same field of view; adjusting a calibration of the video camera based on a measurement of the first light signal; applying the calibration to the second light

signal; measuring a vital sign of the patient from the calibrated second light signal; and outputting the measured vital sign for further processing or display.

**[0009]** In an embodiment, the scale comprises a grey-scale, and the measurement of the first light signal comprises a measurement of a first intensity of at least a portion of the greyscale. In an embodiment, the method includes, at a later time, measuring a second intensity that differs from the first intensity by an amount, and further adjusting the calibration based on the measured second intensity. In an embodiment, the method also includes further adjusting the calibration based on the second intensity comprises adjusting a coefficient in proportion to the amount. In an embodiment, the method includes determining that the amount exceeds a threshold, prior to adjusting the coefficient. In an embodiment, detecting the second light signal comprises combining light from two or more non-contiguous regions of exposed skin of the patient.

**[0010]** In an embodiment, the scale comprises a color map comprising a plurality of colors, and the measurement of the first light signal comprises a measurement of a color value of one of the plurality of colors. In an embodiment, adjusting the calibration comprises comparing the color value to a reference color value and identifying a difference. In an embodiment, the reference color value comprises a baseline color value measured from the first light signal at a first time. In an embodiment, adjusting the calibration comprises determining that the difference exceeds a threshold, and adjusting a coefficient based on the difference.

**[0011]** In an embodiment, the scale comprises a grey-scale, and wherein the measurement of the first light signal comprises a measurement of a white value of the greyscale. In an embodiment, applying the calibration comprises white balancing the first and second light signals.

**[0012]** In an embodiment, the method includes operating a second video camera to monitor a second patient, and adjusting a calibration of the second video camera to match the adjusted calibration of the first video camera.

**[0013]** In an embodiment, the method includes detecting motion based on movement of the scale within the field of view, and generating a motion signal based on the detected motion.

**[0014]** In an embodiment, the method includes measuring oxygen saturation from the calibrated second light signal. The calibrated second light signal comprises two of a red signal, a green signal, and a blue signal, and measuring the oxygen saturation comprises measuring a ratio of the two of the red, green, and blue signals. In an embodiment, detecting the second light signal comprises combining light from two or more non-contiguous regions of exposed skin of the patient.

**[0015]** In an embodiment, the video comprises an optical splitter, and the calibrated second light signal comprises two light signals output from the optical splitter.

**[0016]** In an embodiment, a system for video-based

measurement of a patient's pulse rate includes a video camera positioned remote from a patient, the video camera having a field of view encompassing exposed skin of the patient; a calibration strip positioned within the field of view, the calibration strip comprising a scale viewable by the camera; and a hardware memory coupled to the video camera by wired or wireless communication, the memory storing instructions for instructing a processor to: detect a first light intensity signal from the scale and a second light intensity signal from the exposed skin of the patient; adjust a calibration of the video camera based on a measurement of the first light intensity signal; apply the calibration to the second light intensity signal; measure a pulse rate of the patient from the calibrated second light intensity signal; and output the measured pulse rate for further processing or display. In an embodiment, the calibration strip comprises first and second light emitters. In an embodiment, the calibration strip is devoid of a photodetector.

**[0017]** Independent component analysis may be used, as explained in more detail below. In an embodiment, a method for measuring blood oxygen saturation of a patient includes receiving, from a video camera, a video signal encompassing exposed skin of a patient; extracting from the video signal time-varying red, green, and blue signals; decomposing the red, green, and blue signals into a component signal having a primary frequency at a pulse rate of the patient; identifying, in the component signal, an individual pulse representative of a heart beat; locating a corresponding portion of two of the red, green, and blue signals; and measuring blood oxygen saturation of the patient from the located corresponding portions of the two signals.

**[0018]** In an embodiment, the method includes determining and displaying a pulse rate measured from the primary frequency of the component signal. In an embodiment, an audio beep is triggered in synchrony with the located corresponding portion of one of the two signals or in synchrony with the identified individual pulse in the component signal. In an embodiment, the located portions of the two signals comprise cardiac pulses, and, for each of the two signals, the located cardiac pulse is added to a weighted average pulse. In an embodiment, measuring blood oxygen saturation comprises calculating a ratio of ratios of the weighted average pulses of the two signals. In an embodiment, extracting the red, green, and blue signals comprises selecting pixels within the image frame that exhibit a modulation at the primary frequency. In an embodiment, the selected pixels are non-contiguous.

**[0019]** In an embodiment, extracting the red, green, and blue signals comprises selecting pixels within the image frame exhibiting a modulation that is at the primary frequency and that has an amplitude above a threshold.

**[0020]** In an embodiment, a method for measuring a pulse rate of a patient includes receiving, from a video camera, a video signal having a field of view encompassing exposed skin of a patient; identifying, within the video

signal, regions of pixels that exhibit a modulation above an amplitude threshold; extracting from the identified regions time-varying red, green, and blue signals; decomposing the red, green, and blue signals into a component signal having a primary frequency at a pulse rate of the patient; measuring the pulse rate from the primary frequency of the component signal; and outputting the measured pulse rate for further processing or display.

**[0021]** In an embodiment, the method also includes identifying, in the component signal, individual pulses representative of individual heart beats; for each identified pulse, locating a corresponding portion of two of the red, green, and blue signals; and measuring blood oxygen saturation of the patient from the located corresponding portions of the two signals.

**[0022]** A frequency accumulator may be used, as explained in more detail below. In an embodiment, a method for video-based monitoring of a patient's pulse rate includes generating a video signal from a video camera having a field of view exposed to a patient, the video signal comprising a time-varying intensity signal for each of a plurality of pixels in the field of view; combining the intensity signals within a region of the field of view to produce a regional intensity signal; transforming the regional intensity signal into the frequency domain to produce a regional frequency signal; over a sliding time window, identifying peaks in the regional frequency signal; over a period of time, accumulating the identified peaks; selecting a median frequency from the accumulated peaks; updating a running average pulse rate of a patient, wherein updating comprises: converting the median frequency into a measured pulse rate; and adding the measured pulse rate to the running average to produce an updated average pulse rate; and outputting the updated average pulse rate for display.

**[0023]** In an embodiment, the period of time is one second. In an embodiment, identified peaks from the accumulated peaks are removed based on an age of the identified peaks. In an embodiment, the method includes, at repeated intervals, discarding the accumulated peaks and repeating the steps of identifying peaks, accumulating peaks, selecting the median frequency, updating the running average pulse rate, and outputting the updated average pulse rate.

**[0024]** In an embodiment, adding the measured pulse rate to the running average comprises applying a weight to the measured pulse rate based on a quality of the regional frequency signal. In an embodiment, the quality of the regional frequency signal is measured by a variability of the accumulated peaks over the period of time. In an embodiment, the quality of the regional frequency signal is measured by an amplitude of the accumulated peaks. In an embodiment, the quality of the regional frequency signal is measured by a signal to noise ratio of the regional frequency signal.

**[0025]** In an embodiment, frequency peaks outside of a physiologic limit are discarded. In an embodiment, the measured pulse rate is discarded when it differs from the

average pulse rate by more than a defined amount.

**[0026]** In an embodiment, the method includes updating an average respiration rate of the patient, wherein updating the average respiration rate comprises: selecting a second median frequency from the identified peaks; converting the second median frequency into a measured respiration rate; and adding the measured respiration rate to the average respiration rate to produce an updated average respiration rate; and outputting the updated average respiration rate for display.

**[0027]** In an embodiment, selecting the region of the field of view is based on a strength of modulations of the pixels in the region. In an embodiment, the region comprises two or more non-adjacent groups of pixels.

**[0028]** The frame rate may be adjusted to reject noise, as explained in more detail below. In an embodiment, a method for video-based monitoring of a vital sign of a patient includes receiving a video signal from a video camera having a field of view exposed to a patient, the video signal comprising a time-varying intensity signal for each of a plurality of pixels in the field of view; combining the intensity signals of selected pixels to produce a time-varying regional intensity signal; transforming the regional intensity signal into the frequency domain to produce a regional frequency signal; operating the video camera at a first frame rate during a first period of time, and at a second, different frame rate during a second, subsequent period of time; identifying, in the regional frequency signal, a noise peak at a first frequency during the first period of time that moves to a second, different frequency upon a transition from the first period of time to the second period of time; filtering the regional intensity signal to remove the frequency of the noise peak; and measuring a vital sign of the patient from the filtered regional intensity signal.

**[0029]** In an embodiment, the method also includes identifying a stationary peak that remains stationary in the frequency domain upon the transition from the first period of time to the second period of time, and measuring the vital sign from the identified stationary peak. In an embodiment, the vital sign comprises pulse rate, and measuring the vital sign comprises converting the frequency of the identified stationary peak into the pulse rate. In an embodiment, combining the intensity signals of selected pixels comprises selecting the pixels that exhibit modulations at a shared frequency. In an embodiment, the selected pixels are non-contiguous.

**[0030]** In an embodiment, a method for video-based monitoring of a vital sign of a patient includes receiving a video signal from a video camera having a field of view exposed to a patient, the video signal comprising a time-varying intensity signal for each of a plurality of pixels in the field of view; combining the intensity signals of selected pixels to produce a time-varying regional intensity signal; transforming the regional intensity signal into the frequency domain to produce a regional frequency signal; operating the video camera at a first frame rate during a first period of time, and at a second, different frame

rate during a second, subsequent period of time; identifying, in the regional frequency signal, a stationary peak that is stationary upon a transition from the first period of time to the second period of time; and measuring a vital sign of the patient from the identified stationary peak.

**[0031]** In an embodiment, the method also includes identifying, in the regional frequency signal, a noise peak that is non-stationary upon the transition from the first period of time to the second period of time, and filtering the regional intensity signal to remove the noise peak.

**[0032]** In an embodiment, a method for video-based monitoring of a vital sign of a patient includes receiving a video signal from a video camera having a field of view exposed to a patient, the video signal comprising a time-varying intensity signal for each of a plurality of pixels in the field of view; operating the video camera at a frame rate that changes according to a change trajectory over a period of time; combining the intensity signals of selected pixels to produce a combined intensity signal; transforming the combined intensity signal into the frequency domain to produce a frequency signal; identifying, in the frequency signal, a noise peak that moves in synchrony with the change trajectory; filtering the combined intensity signal to remove the noise peak; and measuring a vital sign of the patient from the filtered intensity signal.

**[0033]** In an embodiment, the method also includes identifying, in the frequency signal, a physiologic peak that is stationary over the period of time, and measuring the vital sign from the physiologic peak. In an embodiment, the physiologic peak corresponds to a physiologic frequency of the patient. In an embodiment, the physiologic frequency is pulse rate. In an embodiment, the physiologic frequency has a period that is smaller than the period of time.

**[0034]** In an embodiment, the change trajectory comprises sweeping the frame rate at a constant sweep rate over the period of time. In an embodiment, identifying the noise peak comprises identifying a peak that moves at the sweep rate. In an embodiment, the change trajectory comprises three or more different, discrete frame rates. In an embodiment, the frame rate is fixed after the noise peak has been identified.

**[0035]** In an embodiment, a method for video-based monitoring of a vital sign of a patient includes receiving a video signal from a video camera having a field of view exposed to a patient, the video signal comprising a time-varying intensity signal for each of a plurality of pixels in the field of view; combining the intensity signals of selected pixels to produce a combined intensity signal; transforming the combined intensity signal into the frequency domain over first and second different time windows to produce first and second frequency transforms; identifying, in the first and second frequency transforms, a noise peak that remains stationary in the first and second frequency transforms; filtering the combined intensity signal to remove the noise peak; and measuring a vital sign of the patient from the filtered intensity signal.

**[0036]** A region of interest may be displayed to a user, as explained in more detail below. In an embodiment, a video-based method for measuring a vital sign of a patient includes receiving a video signal from a video camera having a field of view exposed to a patient; displaying on a display screen a portion of the video signal; receiving a first user input from a user, the first user input identifying a location of a first region of interest within the video signal; extracting from the first region of interest a first intensity signal comprising a time-varying intensity of one or more pixels in the first region; measuring a first vital sign of the patient from the first intensity signal; displaying on the display screen a time-varying modulation of the first intensity signal; receiving a second user input from a user, the second user input indicating that the location has been moved to a second, different region of interest; extracting from the second region of interest a second intensity signal comprising a time-varying intensity of one or more pixels in the second region; measuring a second vital sign of the patient from the second intensity signal; and displaying on the display screen a time-varying modulation of the second color signal.

**[0037]** In an embodiment, the first vital sign comprises heart rate and the second vital sign comprises respiration rate. In an embodiment, the modulations of the first and second color signals are displayed on the display screen simultaneously. In an embodiment, the second user input further comprises a path from the first region of interest to the second region of interest, and wherein the method further comprises identifying a plurality of intermediate regions of interest along the path, extracting an intermediate intensity signal from one of the intermediate regions, and displaying on the display screen a modulation of the intermediate intensity signal. In an embodiment, the first region of interest comprises two or more non-adjacent groups of pixels.

**[0038]** In an embodiment, a video-based method for monitoring a patient includes displaying on a display screen a video signal from a video camera having a field of view exposed to a patient; receiving a first user input from a user, the first user input identifying a location of a region of interest within the video signal; extracting from the region of interest a first intensity signal comprising a time-varying intensity of one or more pixels in the region of interest; displaying on the display screen the first intensity signal; receiving from the user a second user input that moves the region of interest along a path; continually updating the location of the region of interest along the path; continually updating the intensity signal extracted from the moving region of interest; and displaying on the display screen the continually updated intensity signal.

**[0039]** In an embodiment, the method also includes identifying a modulation of the intensity signal, and measuring a physiologic rate of the patient from the modulation. In an embodiment, the physiologic rate is pulse rate. In an embodiment, a transform of the intensity signal into the frequency domain is displayed.

## DETAILED DESCRIPTION

**[0040]** The present invention relates to the field of medical monitoring, and in particular non-contact, video-based monitoring of pulse rate, respiration rate, motion, activity, and oxygen saturation. Systems and methods are described for receiving a video signal in view of a patient, identifying a physiologically relevant area within the video image (such as a patient's forehead or chest), extracting a light intensity signal from the relevant area, filtering those signals to focus on a physiologic component, and measuring a vital sign from the filtered signals. The video signal is detected by a camera that views but does not contact the patient. With appropriate selection and filtering of the video signal detected by the camera, the physiologic contribution to the detected signal can be isolated and measured, producing a useful vital sign measurement without placing a detector in physical contact with the patient. This approach has the potential to improve patient mobility and comfort, along with many other potential advantages discussed below.

**[0041]** As used herein, the term "non-contact" refers to monitors whose measuring device (such as a detector) is not in physical contact with the patient. Examples include cameras, accelerometers mounted on a patient bed without contacting the patient, radar systems viewing the patient, and others. "Video-based" monitoring is a sub-set of non-contact monitoring, employing one or more cameras as the measuring device. In an embodiment, the camera produces an image stack, which is a time-based sequence of images of the camera's field of view. The camera may be considered a "video" camera if the frame rate is fast enough to create a moving, temporal image signal.

**[0042]** Remote sensing of a patient in a video-based monitoring system presents several new challenges. One challenge is presented by motion. The problem can be illustrated with the example of pulse oximetry. Conventional pulse oximetry sensors include two light emitters and a photodetector. The sensor is placed in contact with the patient, such as by clipping or adhering the sensor around a finger, toe, or ear of a patient. The sensor's emitters emit light of two particular wavelengths into the patient's tissue, and the photodetector detects the light after it is reflected or transmitted through the tissue. The detected light signal, called a photoplethysmogram (PPG), modulates with the patient's heartbeat, as each arterial pulse passes through the monitored tissue and affects the amount of light absorbed or scattered. Movement of the patient can interfere with this contact-based oximetry, introducing noise into the PPG signal due to compression of the monitored tissue, disrupted coupling of the sensor to the finger, pooling or movement of blood, exposure to ambient light, and other factors. Modern pulse oximeters employ filtering algorithms to remove noise introduced by motion and to continue to monitor the pulsatile arterial signal.

**[0043]** However, movement in non-contact pulse oxi-

metry creates different complications, due to the extent of movement possible between the patient and the camera, which acts as the detector. Because the camera is remote from the patient, the patient may move toward or away from the camera, creating a moving frame of reference, or may rotate with respect to the camera, effectively morphing the region that is being monitored. Thus, the monitored tissue can change morphology within the image frame over time. This freedom of motion of the monitored tissue with respect to the detector introduces new types of motion noise into the video-based signals.

**[0044]** Another challenge is the contribution of ambient light. In this context, "ambient light" means surrounding light not emitted by components of the medical monitor. In contact-based pulse oximetry, the desired light signal is the reflected and/or transmitted light from the light emitters on the sensor, and ambient light is entirely noise. The ambient light can be filtered, removed, or avoided in order to focus on the desired signal. In contact-based pulse oximetry, contact-based sensors can be mechanically shielded from ambient light, and direct contact between the sensor and the patient also blocks much of the ambient light from reaching the detector. By contrast, in non-contact pulse oximetry, the desired physiologic signal is generated or carried by the ambient light source; thus, the ambient light cannot be entirely filtered, removed, or avoided as noise. Changes in lighting within the room, including overhead lighting, sunlight, television screens, variations in reflected light, and passing shadows from moving objects all contribute to the light signal that reaches the camera. Even subtle motions outside the field of view of the camera can reflect light onto the patient being monitored. Thus new filtering techniques are needed to isolate the physiologic signal from this combined ambient light signal.

**[0045]** If these challenges are addressed, non-contact monitoring such as video-based monitoring can deliver significant benefits. Some video-based monitoring can reduce cost and waste by reducing usage of disposable contact sensors, replacing them with reusable camera systems. Video monitoring may also reduce the spread of infection, by reducing physical contact between caregivers and patients (otherwise incurred when the caregiver places, adjusts, or removes the contact sensor on the patient). Some remote video cameras may improve patient mobility and comfort, by freeing patients from wired tethers or bulky wearable sensors. This untethering may benefit patients who need exercise and movement. In some cases, these systems can also save time for caregivers, who no longer need to reposition, clean, inspect, or replace contact sensors. Another benefit comes from the lack of sensor-off alarms or disruptions. A traditional contact-based system can lose the physiologic signal when the contact sensor moves or shifts on the patient, triggering alarms that are not actually due to a change in physiology. In an embodiment, a video-based system does not drop readings due to sensors moving or falling off the patient (sensor-off) or becoming discon-

nected from the monitor (sensor-disconnect), and thus can reduce nuisance alarms. In an embodiment, a video-based monitor, such as a pulse oximeter, operates without sensor-off or sensor-disconnect alarms. For example, a video-based monitor can trigger an alarm based on stored alarm conditions, where the stored alarm conditions omit a sensor-off or sensor-disconnect alarm.

**[0046]** Various embodiments of the present invention are described below, to address some of these challenges. Figure 2A shows a video-based remote monitoring system 200 and a patient 212, according to an embodiment. The system 200 includes a non-contact detector 210 placed remote from the patient 212. In this embodiment, the detector 210 includes a camera 214, such as a video camera. The camera 214 is remote from the patient, in that it is spaced apart from and does not contact the patient. The camera includes a detector exposed to a field of view 216 that encompasses at least a portion of the patient 212. In some embodiments, the field of view 216 encompasses exposed skin of the patient, in order to detect physiologic signals visible from the skin such as arterial oxygen saturation (SpO<sub>2</sub> or SvidO<sub>2</sub>). The camera generates a sequence of images over time. A measure of the amount, color, or brightness of light within all or a portion of the image over time is referred to as a light intensity signal. In an embodiment, each image includes a two-dimensional array or grid of pixels, and each pixel includes three color components - for example, red, green, and blue. A measure of one or more color components of one or more pixels over time is referred to as a "pixel signal," which is a type of light intensity signal. The camera operates at a frame rate, which is the number of image frames taken per second (or other time period). Example frame rates include 20, 30, 40, 50, or 60 frames per second, greater than 60 frames per second, or other values between those. Frame rates of 20-30 frames per second produce useful signals, though frame rates above 50 or 60 frames per second are helpful in avoiding aliasing with light flicker (for artificial lights having frequencies around 50 or 60 Hz).

**[0047]** The detected images are sent to a monitor 224, which may be integrated with the camera 214 or separate from it and coupled via wired or wireless communication with the camera (such as wireless communication 220 shown in Figure 2A). The monitor 224 includes a processor 218, a display 222, and hardware memory 226 for storing software and computer instructions. Sequential image frames of the patient are recorded by the video camera 214 and sent to the processor 218 for analysis. The display 222 may be remote from the monitor 224, such as a video screen positioned separately from the processor and memory.

**[0048]** Figure 2B shows the system 200 being implemented to monitor multiple patients, such as patients 212A and 212B. Because the detector 214 in the system is non-contact, it can be used to monitor more than one patient at the same time. A method for this implementation will be described in further detail below.

**[0049]** Two example image frames 300A and 300B are shown in Figures 3A and 3B, respectively. In an embodiment, these image frames are recorded by the system 200. Each image frame includes a patient's head 312 and upper torso 310 in the field of view. The processor has identified a head region 314 within each image frame 300A, 300B. The head region 314 includes at least a portion of the patient's head, such as the face. In some embodiments, the processor also infers a chest region 316, based on the size and location of the head region 314 and empirical ratios of head and chest sizes and shapes. For example, from a rectangular face region of width  $w$  and height  $h$ , a forehead region may be inferred of a size  $0.7*w$  and  $0.3*h$ , centered horizontally and positioned with its top edge moved down from the top of the face region by a distance  $0.25*h$ . From the same rectangular face region, a chest region may also be inferred at a size of  $2*w$  and  $0.75*h$ , centered horizontally and positioned with its top edge below the bottom of the face region by a distance  $0.25*h$ .

**[0050]** In an embodiment, the video camera records multiple sequential image frames (such as image frames 300A and 300B) that each include the head region 314 and chest region 316. The pixels or detected regions in these sequential images exhibit subtle modulations caused by the patient's physiology, such as heartbeats and breaths. In particular, the color components of the pixels vary between the frames based on the patient's physiology. In one embodiment, the camera employs the Red/Green/Blue color space and records three values for each pixel in the image frame, one value each for the Red component of the pixel, the Blue component, and the Green component. Each pixel is recorded in memory as these three values, which may be integer numbers (typically ranging from 0 to 255 for 8-bit color depth, or from 0 to 4095 for 12-bit color depth) or fractions (such as between 0 and 1). Thus, three one-dimensional vectors for each pixel in the field of view can be extracted from the video signal.

**[0051]** These Red, Green, and Blue values change over time due to the patient's physiology, though the changes may be too subtle to be noticed by the naked human eye viewing the video stream. For example, the patient's heartbeat causes blood to pulse through the tissue under the skin, which causes the color of the skin to change slightly - causing the value corresponding to the Red, Green, or Blue component of each pixel to go up and down. These changes in the pixel signals can be extracted by the processor. The regions within the field of view where these changes are largest can be identified and isolated to focus on the physiologic signal. For example, in many patients, the forehead is well-perfused with arterial blood, so pixels within the patient's forehead exhibit heartbeat-induced modulations that can be measured to determine the patient's heartrate.

**[0052]** To focus on this physiologic signal, the processor identifies a region of interest (ROI) within the image frame. In an embodiment, the region of interest includes

exposed skin of the patient, such that the physiologic properties of the skin can be observed and measured. For example, in the embodiment of Figure 3A, one region of interest includes a forehead region 330, which includes part of the patient's forehead. The processor determines the location of the patient's forehead within the head region 314, for example based on empirical ratios for a human face, and divides the forehead into distinct regions, for example, regions 1A, 2A, and 3A. In another embodiment, the region of interest does not include exposed skin. For example, in Figure 3A, another region of interest includes the chest region 316 (which may be covered by clothing, bedding, or other materials on the patient). Pixels in this region may fluctuate with the patient's respiration rate, enabling that rate to be measured even without viewing exposed skin of the patient.

**[0053]** Within an individual region of interest, the Red components of the pixels in that region are combined together to produce one time-varying Red pixel signal from that region. The same is done for the Blue and Green pixels. The result is three time-varying pixel signals from each region, and these are plotted in Figure 4A. The plots in Figure 4A are derived from the regions 1A, 2A, 3A, and 316 of Figure 3A. Figure 4A also shows a plot labeled "Combined Forehead." The Combined Forehead plot shows the combined pixel signals from all three identified regions 1A, 2A, and 3A, meaning that the Red components from all three regions are combined together and plotted over time, as are the Green components and the Blue components. Different sub-sets of regions can be combined together to produce different combinations of pixel signals. Though three forehead regions 1A, 2A, and 3A are shown in Figure 3A, the forehead, or any other area of interest, can be sub-divided into more or fewer regions, in various shapes or configurations. (Other examples are described in more detail below with reference to Figures 5A and 5B.) Pixel signals can be combined by summing or averaging or weighted averaging. In an embodiment, the combined pixel signals are obtained by averaging the Red (or Blue, or Green) color values of the pixels within the region, so that regions of different sizes can be compared against each other.

**[0054]** The pixels within a region may be combined together with a weighted average. For example, within a region, some pixels may exhibit stronger modulations than other pixels, and those stronger-modulating pixels can be weighted more heavily in the combined pixel signal. A weight can be applied to all of the pixels that are combined together, and the weight can be based on quality metrics applied to the modulating intensity signal of each pixel, such as the signal to noise ratio of the intensity signal, a skew metric, an amplitude of a desired modulation (such as modulations at the heart rate or respiration rate), or other measurements of the signal. Further, some pixels within the region may be chosen to be added to the combined pixel signal for that region, and other pixels may be discarded. The chosen pixels need not be adjacent or connected to each other; disparate pixels can be

chosen and combined together to create the resulting signal.

**[0055]** The plots in Figure 4A show a clear pattern of repeating modulations or pulses over time. The pulses in each region 1A, 2A, 3A and in the Combined Forehead plot are caused by the patient's heart beats, which move blood through those regions in the patient's forehead, causing the pixels to change color with each beat. The heart rate of the patient can be measured from these signals by measuring the frequency of the modulations. This measurement can be taken via a frequency transform of the signal (discussed below with reference to Figure 10A and Figure 4B) or via a pulse recognition algorithm that identifies each pulse in the signal (for example, by pulse size and shape, by zero crossings, maximums, or minimums in the derivative of the signal, and/or by checking the skew of the derivative of the signal to identify a pulse as a cardiac pulse, which has a characteristically negative skew). The modulations in the plot of the Chest region, in Figure 4A, are caused by the patient's breaths, which cause the chest to move in correspondence with the breathing rate. The patient's breathing/respiration rate can be measured from this signal in the same way as just described for the heart rate (except for the skew approach). Respiration rate can be identified from a region of the patient that moves with each breath, such as the chest, but need not include exposed skin.

**[0056]** Figure 4B shows plots of the pixel streams from the corresponding regions in Figure 3B. However, in this case, the individual Red, Green, and Blue values within each region have been combined together, such as by summing or averaging, to produce one time-varying signal from each region instead of three separate Red, Green, and Blue signals. By viewing one combined signal from each region, the frequency of the heart rate or respiration rate may emerge more clearly. Figure 4B also shows a Fast Fourier Transform (FFT) in the Chest Region plot. The FFT identifies the frequency content of the Chest signal, which reveals a primary frequency peak and harmonics. The primary frequency peak is the patient's respiration rate.

**[0057]** Though many embodiments herein are described with reference to pixels and pixel values, this is just one example of a detected light intensity signal. The light intensity signals that are detected, measured, or analyzed may be collected from larger regions or areas, without differentiating down to groups of pixels or individual pixels. Light signals may be collected from regions or areas within an image, whether or not such regions or areas are formed from pixels or mapped to a spatial grid. For example, time-varying light signals may be obtained from any detector, such as a camera or light meter, that detects a unit of light measurement over time. Such units of light measurement may come from individual pixels, from groups or clusters of pixels, regions, sub-regions, or other areas within a field of view. It should also be noted that the term "pixel" includes larger pixels that are themselves formed from aggregates, groups, or clusters

of individual pixels.

**[0058]** In an embodiment, the Red, Green, and Blue values from the camera are converted into different color spaces, and the color space that provides the largest or most identifiable physiologic modulations is chosen. In an embodiment, color values are converted into a combination of a color value and a separate brightness value, so that changes in room brightness can be analyzed independently of color or hue. Alternative color spaces (such as YCrCb, CIE Lab, CIE Luv) can separate light intensity from chromatic changes better than the RGB color space. Processing the chromatic component in those spaces can reveal physiological modulation better than in RGB space, when overall scene light intensity is changing. Assessing pixel signals based on chromatic channels in these spaces can increase the robustness of the algorithm and/or increase the range of conditions in which physiological signal extraction is possible. Though the Red/Green/Blue color scheme is often presented here in the examples, it should be understood that other color schemes or color spaces can be utilized by these systems and methods.

**[0059]** Figures 3A and 3B depict five regions of interest - three squares in the forehead, the combination of all three squares together, and one rectangular chest region. In other embodiments, regions of interest can have various shapes, configurations, or combinations. Examples are shown in Figures 5A and 5B. In the embodiment of Figure 5A, a monitor 500A displays an image frame 502, which depicts a region of interest on a patient, in this case a face region, and in particular a forehead region 503. The face region is further divided into a grid 504, segmenting the face region into smaller individual regions. Within this grid 504, individual regions of interest 506A, 506B, 506C, 506D, ...506N are identified. The regions of interest 506A-N are regions that include pixels or detected areas that exhibit a physiologic characteristic of the patient. A sub-set of the regions of interest can be chosen to focus on a particular physiologic characteristic that is reflected in the pixels in those regions.

**[0060]** In one embodiment, the selected regions of interest (for measurement of a vital sign) are completely enclosed within the patient region, such as the face or a smaller area such as the forehead. For example, in Figure 5A, the regions 506A-C are completely contained within the patient's forehead region 503. No portion of regions 506A-C includes pixels outside of the patient's forehead. These regions 506A-C are used to identify a physiologic signal and calculate a vital sign, such as the patient's heartrate, as described above. By enclosing the regions within a physiological area, such as the forehead, according to some embodiments, the signal to noise ratio of the desired physiologic signal increases.

**[0061]** In another embodiment, the selected regions of interest may be non-adjacent to each other, or non-contiguous. For example, in Figure 5A, non-adjacent regions 506A and 506D may both include pixels that exhibit large modulations correlated with the patient's heartrate, as

compared to the other regions. Regions located over large arteries may exhibit larger modulations with heartrate than other regions, for example. In an embodiment, the intensity signals from regions 506A and 506D are averaged together to create a combined signal, and the heartrate measured from that combined signal. Different non-adjacent regions may be chosen for other vital signs, such as respiration rate or oxygen saturation. In an embodiment, heart rate and oxygen saturation are calculated from a combined signal from a first group of non-adjacent pixels or regions, and respiration rate is calculated from a different combined signal from a second, different group of non-adjacent pixels or regions.

**[0062]** In an embodiment, regions of interest within the image frame are selected based on the modulations exhibited by the pixels in each region. Within an image frame, a sub-set of regions may be first identified as candidate regions for further processing. For example, within an image frame, an area of exposed skin of a patient is identified by facial recognition, deduction of a forehead region, user input, and/or skin tone detection. These areas are identified as the regions of interest for further processing. In an embodiment, facial recognition is based on Haar-like features (employing a technique that sums pixel intensities in various regions and differences between sums). A method includes identifying these regions of interest, extracting pixel signals from each region, quantifying the magnitude of physiological modulations exhibited by each pixel signal, selecting regions with strong modulations (such as modulations with an amplitude above a threshold), combining the selected pixel signals together (such as by averaging), and measuring a vital sign from the combined signal. In an embodiment, all sub-regions (such as grids) in the image (or a portion of the image, such as a patient region) are processed, and grid cells that exhibit coherent pulsatile components are combined to generate the pixel signals from which the physiologic measurements are taken.

**[0063]** Selecting non-adjacent regions enables the system to focus on the pixels or regions that carry the physiologic signal with the highest signal to noise ratio, ignoring other areas in the image frame that are contributing a relatively higher degree of noise, such as pixels that do not vary much with heart rate, but that might vary due to a passing shadow or patient movement. The system can focus on pixels that represent the desired vital sign, thereby increasing the signal-to-noise ratio (SNR) of the analyzed signal. With signals from several regions available, the signals with the strongest SNR can be chosen, and signals with weak SNR can be discarded. The chosen signals can be combined together to produce a signal with a strong physiologic component.

**[0064]** Referring to Figure 5A, the size of the cells within the grid 504 can affect the computation of the resulting pixel signals. If the cells in the grid are very small (such as 10 pixels by 10 pixels), the number of cells increases, causing the number of computations and available signals to increase. The variability of the signals also in-

creases with very small cell sizes. For example, a passing shadow or a twitch can affect a very small area of skin. If a region of interest is wholly contained within that affected area, the signal from that region will become noisy. Larger regions provide a degree of spatial smoothing that reduces susceptibility to such noise, but regions that are too large in size may obscure the physiologic signal. An example of a region of a good size for processing a physiologic signal is approximately one square centimeter (though more or less may also be useful - for example a whole forehead may be used, or an individual pixel). If far away from the subject, a camera may use less pixels. The selection of region size also depends on the resolution of the image, which may depend on the available hardware. Moreover, resolution and frame rate may be inter-related, in that increasing resolution may reduce frame rate. A compromise is necessary between high enough resolution to capture the modulating pixels, and a fast enough frame rate to track those modulations over time. Frame rates over 10 Hz are sufficient for cardiac pulses, and over 2-3 Hz for respiration modulations. Frame rates above about 50 or 60 frames per second are generally less subject to aliasing frequencies introduced by artificial lighting. Sampling from a few hundred pixels (such as over 200 or over 300 pixels) has been sufficient to isolate a physiologic modulation above ambient noise.

**[0065]** The selected regions of interest can change over time due to changing physiology, changing noise conditions, or patient movement. In each of these situations, criteria can be applied for selecting a pixel, group of pixels, or region into the combined signal. Criteria are applied to enhance the physiologic signals by reducing or rejecting contributions from stationary or non-stationary non-physiologic signals. Criteria can include a minimum SNR, a minimum amplitude of physiologic modulations, a minimum variability of the frequency of modulations (to reject non-physiologic, static frequencies), a skew metric (such as modulations that exhibit a negative skew), pixels with values above a threshold (in the applicable Red, Green, or Blue channel), pixels that are not saturated, or combinations of these criteria. These criteria can be continually applied to the visible pixels and regions to select the pixels that meet the criteria. Some hysteresis may be applied so that regions or pixels are not added and removed with too much chatter. For example, pixels or regions must meet the criteria for a minimum amount of time before being added to the combined signal, and must fail the criteria for a minimum amount of time before being dropped. In another example, the criteria for adding a pixel or region to the combined signal may be stricter than the criteria for removing the pixel or region from the combined signal.

**[0066]** For example, in an example involving motion, when the patient turns his or her head, the regions of interest that previously demonstrated heart rate with the best amplitude are no longer visible to the camera, or may be covered in shadow or over-exposed in light. New

regions of interest become visible within the field of view of the camera, and these regions are evaluated with the criteria to identify the best candidates for the desired vital sign. For example, referring to Figure 5A, cells or groups of pixels at the edges of the forehead region 503 can be added or removed from the combined signal during motion as they enter and exit the forehead region. This method enables the monitoring system to continue to track the vital sign through movement of the patient, even as the patient moves or rotates with respect to the camera.

**[0067]** Selected regions may also change over time due to changing physiology. For example, these regions can be updated continually or periodically to remove pixels that do not satisfy the criteria for vital sign measurement, and add new pixels that do satisfy the criteria. For example, as the patient's physiology changes over time, one region of the forehead may become better perfused, and the pixels in that region may exhibit a stronger cardiac modulation. Those pixels can be added to the combined light signal to calculate the heart rate. Another region may become less perfused, or changing light conditions may favor some regions over others. These changes can be taken into account by adding and removing pixels to the combined signal, to continue tracking the vital sign.

**[0068]** Selected regions may also change over time due to changing noise conditions. By applying the criteria over time, pixels or regions that become noisy are removed from the combined light intensity signal, so that the physiologic signal can continue to be monitored via pixels or groups that are less noisy. These updates can be made continually.

**[0069]** In another embodiment, as shown in Figure 5B, individual pixels 508A-N within the image frame 502, rather than regions or groups of contiguous pixels, are selected and summed together to produce a signal from which a patient vital sign can be measured. In Figure 5B, the patient region need not be divided into sub-regions, such as the grid 504 shown in Figure 5A. Rather, individual pixels 508 within the patient region are evaluated, and the pixels that modulate in correlation with the desired vital sign are selected and summed together. These pixels need not be adjacent or in a near vicinity of each other.

**[0070]** Figure 5E shows a method for video-based monitoring of a patient's vital signs, according to an embodiment. The method includes receiving a video signal from a video camera at 511. The video signal includes a plurality of sequential image frames, each image frame having a field of view that includes exposed skin of a patient, such as the face or forehead. The method includes segmenting a first image frame into a plurality of regions at 512, and then, for each region, extracting from the video signal a time-varying color signal at 513. In an example, three time-varying color signals are extracted from each region, corresponding to red, green, and blue pixel values. The method includes identifying a frequency content of each color signal at 514, and selecting regions

that have a shared frequency content at 515. The shared frequency content is a modulation at a shared frequency. For example, two regions that both exhibit color signals that modulate at the patient's heart rate, such as a frequency of 60 beats per minute, are selected. In an embodiment, the shared modulation must pass criteria, such as those described above, to select the desired regions. For example, an amplitude threshold for the modulation frequency can be applied as a criterion for selecting regions. In an embodiment, the regions that satisfy this criterion are non-adjacent to each other; they do not need to be in contact with each other or next to each other on the patient. Rather, regions that exhibit a shared modulation at a physiologic frequency, above a noise threshold, are selected even if they are located at disparate, non-contiguous locations across the patient.

**[0071]** Once the desired regions are selected, the method includes combining the color signals of the selected regions at 516, and measuring a vital sign from the combined color signal at 517, such as measuring heart rate from the identified frequency. The vital sign is output for further processing or display at 518. The calculated vital sign can be added to a long-term running average, or a weighted average, where the weight is based on quality metrics such as signal to noise ratio or vital sign variability.

**[0072]** The combined light signal can be used to calculate statistics, such as an amplitude of the physiologic frequency (in the time or frequency domain), a variability of the frequency over time, a variability of the intensity or color of the selected pixels over time, a skew of the modulations, or a signal to noise ratio. Skew is a useful metric because cardiac pulses tend to have a negative skew. Thus, modulations of pixels that exhibit a negative skew may be more likely to be physiologic. In an embodiment, one or more statistics are calculated, and then used to apply a weight to each color signal (from an individual pixel or from a region) that is being combined. This method results in a weighted average that applies more weight to the pixels that exhibit modulations that are stronger or more likely to be physiologic. For example, pixels that modulate with a strongly negative skew, or a high signal to noise ratio, can be weighted more heavily. The criteria used to select regions can also be used to assign weights; for example, regions or pixels that meet a first, stricter set of criteria may be combined with a first, higher weight, and regions or pixels that meet a second, looser set of criteria may be combined with a second, lower weight.

**[0073]** In an embodiment, a weight can also be applied to the vital sign that is calculated from the combined light signal. Each time the vital sign is calculated, a weight can be determined based on current quality measures or statistics from the combined light signal. The newly calculated vital sign is then added to a longer-term running average, based on the weight. For example, the patient's heart rate can be calculated from the combined light signal once per second. An associated weight can be calculated based on the criteria applied to the com-

bined light signal. The weight is reduced when statistics indicate that the light signal may be unreliable (for example, the amplitude of the modulations drops, or the frequency becomes unstable, or the intensity changes suddenly) and increased when statistics indicate that the light signal is reliable.

**[0074]** Furthermore, different combinations of pixels (and/or regions) may be selected for different vital signs of the patient. For example, a first group of pixels and/or regions is summed together to produce a signal that modulates with heart rate, and a second group of pixels and/or regions is summed together to produce a signal that modulates with respiration rate. This approach is demonstrated in Figures 5C and 5D, which each show a light intensity signal over the same span of time from the same video signal for the same patient, from different regions, such as groups of pixels. The pixels chosen for the plot in Figure 5C exhibit relatively large fluctuations correlated with the patient's respiration. This is shown by the large baseline modulations 520, with period P1, in the plotted pixel signal. The frequency of the modulations 520 is the patient's respiration rate, such as 5-20 breaths per minute. By contrast, the pixels chosen for the plot in Figure 5D do not fluctuate as dramatically with the patient's respiration, but they do fluctuate with the patient's heart rate, as shown by the modulations 530 with shorter period P2. The frequency of these modulations is the patient's heart rate, such as 40-200 beats per minute. These two different plots shown in Figures 5C and 5D reflect different vital signs of the patient, based on the same video stream from the same camera taken over a single period of time. By creating combined pixel signals from appropriately selected pixels or regions, various physiologic signals emerge from the video images.

**[0075]** Accordingly, in an embodiment, a method is provided for measuring different vital signs from different regions. These groups can include individual pixels, disparate pixels, contiguous regions, non-contiguous regions, and combinations of these. Pixels combined into one group exhibit a common modulation, such as a frequency of modulation of color or intensity. For example, heart rate can be measured from the frequency of modulation of a first group of pixels, and respiration rate can be measured from the frequency of modulation of a second group of pixels. Oxygen saturation can be measured from either group; in one embodiment, oxygen saturation is measured from the pixels that show strong modulation with heart rate. Specifically, oxygen saturation is measured as a ratio of ratios of the cardiac pulsatile components of two of the signals (such as Red and Green, or Red and Blue) (as described in more detail below).

**[0076]** In an embodiment, a user can view a video image, specify a region of interest, and drag and drop the region across the video image to view changes in modulations in real-time. For example, referring to Figure 5B, a monitor 500B displays a video image 502 that accepts inputs from a user. A user can use mouse pointer 509 (or other input) to highlight a first area 507A, and view

the resulting pixel signals such as the signal shown in Figure 5C and vital signs measured from that signal. The user can then drag and drop the area of interest to a second area 507B and view the resulting signal and vital signs, such as the signal shown in Figure 5D. In this way, the user can view in realtime how the modulations of the signal change based on the selected area of interest. In area 507A, the video signal shows strong respiration modulations (see Figure 5C), while in area 507B, the video signal shows strong cardiac modulations (see Figure 5D). The user can view the video signal in real-time as it moves along the path from 507A to 507B, to see how the modulations change as the region of interest moves. The user can also view the pixel signals shown in Figures 5C and 5D at the same time, to evaluate different vital signs from different regions of interest, at the same time.

**[0077]** A method for monitoring a patient by viewing these different modulations across different regions of interest is outlined in Figure 5F. The method includes displaying a video signal at 521, and receiving a first user input identifying a region of interest within the video image at 522. The method includes extracting a color signal from the region of interest at 523, and displaying the color signal at 524. The method then includes receiving a second user input that moves the region of interest along a path (such as from 507A to 507B in Fig. 5B) at 525. The method includes continually updating the location of the region of interest in accordance with the second user input at 526, continually updating the color signal from the region of interest at 527, and displaying the updated color signal at 528. This enables a user to dynamically change the region of interest and view the resulting extracted video signal, to dynamically see the modulations at any point in the field of view. In addition to displaying the color signal, vital signs can be calculated from the moving region of interest and displayed to the user.

**[0078]** In an embodiment, a video-based method for measuring a vital sign of a patient includes receiving a video signal, displaying on a display screen an image frame from the video signal, and receiving from a user a first user input that identifies a location of a first region of interest within the image frame. The method includes extracting from the first region of interest a first color signal comprising a time-varying intensity of one or more pixels in the first region, detecting a modulation of the first color signal, and measuring a first vital sign of the patient from the modulation of the first color signal. The first vital sign and/or the modulation may be displayed. The method also includes receiving a second user input indicating that the location has been moved to a second, different region of interest. The method then includes extracting from the second region of interest a second color signal, detecting a modulation of the second color signal, measuring a second vital sign of the patient from the modulation of the second color signal, and displaying the modulation and/or second vital sign. In an embodiment, the method includes identifying a plurality of intermediate

regions of interest along the path from the first to the second region of interest, extracting an intermediate color signal from one of the intermediate regions, and displaying on the display screen a modulation of the intermediate color signal.

**[0079]** In an embodiment, the desired pixels are chosen based on a ratio of modulations in the pixel signals. A ratio  $R$  is defined as  $A_H / A_R$ , where  $A_H$  is the cardiac pulse amplitude, and  $A_R$  is the respiratory modulation amplitude. The region where  $R$  is maximum (or above a suitable threshold) can be used to determine heart rate, and the region where  $R$  is minimum (or below a suitable threshold) can be used to determine respiratory rate. A method according to this embodiment is shown in Figure 6A. Regions may be increased or decreased in size, or discrete regions or pixels combined together, to obtain a combined pixel signal with an optimal or desired ratio  $R$ .

**[0080]** As discussed above, a region of interest can be formed based on pixels that modulate with the patient's heart rate. Heart rate can then be calculated from the frequency content of that pixel signal. An example method for calculating heart rate is shown in Figure 10B. The method includes capturing video, acquiring and averaging color signals (shown as  $p_R$ ,  $p_G$ , and  $p_B$  for "photoplethysmogram" red, green, and blue) within a well-perfused ROI, de-noising the signal, performing an FFT (fast Fourier transform) operation over a sliding time window (such as 20 seconds) to identify frequency components of the signals, finding peak frequencies, and accumulating peaks over a period of time (such as one second). De-noising includes filtering the signal to remove noise sources and frequencies outside of a known physiologic range. Examples of filtering operations to remove noise are described below with reference to Figures 17A and 17B. In accumulating peaks, the method may add frequencies multiple times based on their relative height, and may add harmonics of already-added frequencies only once. Frequency peaks are added to the accumulator at the frame rate, such as 25-30 times per second.

**[0081]** Then, once per second, the method finds a median frequency from the accumulated ones, and determines heart rate from the median frequency. The determined heart rate is added to an ongoing average, and then posted for display. As noted in the figure, different averaging techniques may be employed for the externally-posted heart rate as well as for an internally-maintained running average, such as to apply additional smoothing to the externally-posted heart rate. When multiple peaks are present, additional filtering can be applied to determine the most likely heart rate. For example, frequency peaks outside of known physiologic limits for heart rate (such as below 40 or above 250 beats per minute) are rejected. Knowledge of the patient's previous heart rate is also useful, as the heart rate is unlikely to jump a large amount (such as 2.5% of the current heart rate, or another percentage, or a value such as 15 or 20 beats per minute) within 1 second, so such frequency peaks can be rejected as noise. Within the acceptable

peaks, the strongest peak is selected as the patient's heart rate. When the median frequency is rejected as noise, the previously-calculated heart rate is held for one additional second, while the next group of peaks is accumulated, and the range for an acceptable heart rate is increased. When the new group of peaks is assessed, a median frequency is picked, and a new heart rate calculated, the acceptable range is re-set to its normal size, around the new average heart rate. The same method can be used to determine respiration rate, within different frequency ranges and time windows, applied to the same or different pixel signals.

**[0082]** Figure 10B also includes a cross-correlation process that cross-correlates the frequency spectrums of the three color signals to amplify the results. All four resulting spectrums are analyzed to select and accumulate peaks. A cross correlated spectrum can be calculated by multiplying or summing existing spectrum together. An individual spectrum can be scaled before being combined based on signal quality. For example, because most RGB cameras have twice the number of green pixels compare to red and blue ones, the Green signal is usually better and can be weighted above Red and Blue. This method can follow the strongest peaks around the spectrum over time, as the patient's physiology (such as respiration rate and heart rate) changes.

**[0083]** In an embodiment, a method for monitoring a patient's heart rate includes generating a video signal from a video camera having a field of view encompassing exposed skin of a patient. The video signal includes a time-varying intensity signal for each of a plurality of pixels in the field of view. The method includes combining the intensity signals within a region of the field of view to produce a regional intensity signal, and transforming the regional intensity signal into the frequency domain to produce a regional frequency signal. The region may be selected based on a strength of modulations of intensity signals in the region. The region may include non-adjacent areas or pixels. Over a sliding time window, peaks in the regional frequency signal are identified, and then over a period of time (such as one second), the identified peaks are accumulated. The method includes selecting a median frequency from the identified peaks, and updating a running average heart rate of a patient, which includes converting the median frequency into a measured heart rate and adding the measured heart rate to the running average. The updated average heart rate is output for display. The method may also include removing identified peaks from the accumulated peaks when they reach an age limit. The method may also include discarding frequency peaks outside of a physiologic limit, or discarding the measured heart rate when it differs from the average heart rate by more than a defined amount. The method may also include discarding frequency peaks if they are sub-harmonics of already identified peaks.

**[0084]** An example frequency transform of a pixel signal from a region of interest is shown in Figure 10A. This

figure shows three (Red, Green, and Blue) pixel signals over time and the FFT operation, which is applied to a 20-second window of the cross-correlated spectrum of all three signals. The FFT shows a strong peak at 66.0 beats per minute. In the method of Figure 10B, these peaks are added to the frequency accumulator, the median peak is identified, and the patient's heart rate calculated from the median peak.

**[0085]** The non-contact video monitoring system provides many benefits over traditional contact sensors, and also enables monitoring in new and difficult situations. In one example, the non-contact video-based monitoring system can be used to measure vital signs in patients who are not able to tolerate a contact-based sensor, such as patients with skin trauma. These patients could include burn victims, or patients with other sensitive skin conditions. In another example, the non-contact video-based monitoring system can be used to measure multiple patients at the same time (see Figure 2B). A method for monitoring two or more patients at the same time includes orienting the field of view of the camera to encompass two or more patients. In an embodiment, the camera is oriented such that the field of view encompasses exposed skin of each patient, and groups of pixels that exhibit physiologic modulations are identified for each respective patient. A single camera system can then be used to measure vital signs from multiple patients, such as patients on a general care floor, or to track movement of patients within a room or ward.

**[0086]** The vital signs measured from the video signal can be used to trigger alarms based on physiologic limits (for example, high or low heart rate, SpO<sub>2</sub>, or respiration rate alarms). The video signals, the measured vital signs, and triggered alarms can be used by clinicians to identify patients in distress, provide clinical intervention, apply a treatment, support a diagnosis, or recommend further monitoring. The vital signs measured from the video signals may be further processed to arrive at a final value that can be displayed or compared to alarm limits. Further processing may include adding the vital sign to a running average (such as an infinite impulse response filter) to smooth out variability, rejecting outlier vital sign measurements that are not supported by known physiological limits (such as a newly calculated heart rate that varies by more than a physiologically expected amount, as discussed above), increasing or decreasing a weight applied to the vital sign, calculating statistics relating to the vital sign, or other processing steps. The result is a final number, derived from the vital sign measurement from the intensity signal, and this final derived number can be displayed, stored, or compared to alarm limits.

#### Oxygen Saturation

**[0087]** According to an embodiment of the invention, the Red/Green/Blue pixel streams from identified areas of the patient's exposed skin can be used to determine arterial oxygen saturation (SpO<sub>2</sub>). Traditional pulse oxim-

eters employ contact-based sensors, which include two emitters (typically light emitting diodes, LED's) and a photodetector. The emitters are positioned on the sensor to emit light directly into the patient's skin. The emitters are driven sequentially, so that light of each wavelength can be separately detected at the photodetector, resulting in two time-varying light intensity signals. The wavelengths are chosen based on their relative absorption by oxygenated hemoglobin in the blood. Typically one wavelength falls in the red spectrum and the other in infrared. The patient's arterial oxygen saturation can be measured by taking a ratio of ratios (ROR) of the two signals - that is, by taking a ratio of the alternating component (AC) of each signal to its direct, non-alternating component (DC) and dividing the red ratio by the infrared ratio.

**[0088]** In a video-based system, the Red/Green/Blue pixels or regions detected by the camera provide three light intensity signals that potentially can be used in a similar ratio of ratios calculation, such as by dividing the ratios of any two of the three signals. However, many standard video cameras do not detect light in the infrared wavelengths. Moreover, for many video cameras, the wavelengths of light detected in each of the Red, Green, and Blue components overlap. For example, the video camera 214 (see Figure 2A) may include an image sensor with broad spectrum red, green, and blue detectors. The wavelengths detected by these detectors overlap, and are not chosen specifically for their relative absorption by oxygenated hemoglobin. As a result, measuring a ratio of ratios from two of the three signals does not provide an absolute, calibrated SpO<sub>2</sub> value. However, such a ratio of ratios can be used to track the trend of the patient's actual SpO<sub>2</sub> value.

**[0089]** Such a trend is shown in Figure 6B. The top plot in Figure 6B shows an SpO<sub>2</sub> value from a calibrated, contact-based pulse oximeter. It also shows two heart rate signals, one taken from the same pulse oximeter and the other from a video signal. It is readily apparent that the video-based heart rate signal tracks the oximeter-based heart rate signal very closely, providing good absolute correlation.

**[0090]** The bottom plot in Figure 6B shows three different SpO<sub>2</sub> values from a video signal, one for each pair of signals. The top trace is from a ratio of ratios calculation of the Red and Green signals, the middle is the Red and Blue signals, and the bottom is the Green and Blue signals. These three traces can be compared with the calibrated SpO<sub>2</sub> value plotted above, from the conventional contact pulse oximeter. It is clear from Figure 6B that all three traces correlate with the calibrated SpO<sub>2</sub> plot, in that they trend up or down in proportion to the calibrated SpO<sub>2</sub> plot. However the absolute values (shown in the y-axes in Figure 6B) of the video-based SpO<sub>2</sub> traces do not match the calibrated SpO<sub>2</sub> value itself. The calibration of the SvidO<sub>2</sub> against SpO<sub>2</sub> may be performed by linear regression, whereby the coefficients of the regression model are applied to the SvidO<sub>2</sub> to estimate the

absolute SpO2 values.

**[0091]** In an embodiment, the video-based SpO2 measurement is used as a trend indicator, rather than as a measurement of an accurate SpO2 numerical value. For example, it is apparent from the Blue-Red trace that the SpO2 value remains stable until time t1, begins to change at time t1, decreases until time t2, remains stable at low oxygenation until time t3, increases again until time t4, and thereafter remains stable again. The Blue-Red trace can thus be used as a trend indicator, to provide an alert that the patient's SpO2 value is changing, and can even indicate whether the SpO2 value is increasing or decreasing, and an indication of the rate of increase or decrease. This information can be used to provide an early warning to a caregiver that the patient needs attention, such as by attaching a traditional contact-based pulse oximeter to obtain a numerically accurate reading of the patient's SpO2 value which can be used to determine a diagnosis or treatment.

**[0092]** In another embodiment, the SpO2 value measured from a pair of the Red/Green/Blue pixel streams is calibrated to an accurate numerical value. Calibration can be done by comparing the video-based SpO2 value to the value from a reference contact-based oximeter, to identify an offset between them. This offset is used to determine a scaling factor that is applied to the ROR calculation from the video signal. For example, the scaling factor can be a coefficient multiplied to the video ROR, or an offset added or subtracted from the video SpO2, or both. This offset and/or coefficient can be used until the next recalibration. Recalibration can be done when a set time has expired, or when the video SpO2 trend shows a marked change in SpO2.

**[0093]** Figure 7 shows a method of calibrating a video-based SpO2 measurement, according to an embodiment of the invention. The method includes performing a spot check with a contact oximeter at 701, comparing the oximeter SpO2 to the video SpO2 (also called  $S_{vid}O_2$ ) at 702, and determining the calibration between the two values (such as an offset, scaling factor, and/or coefficient) at 703. The method then includes measuring SpO2 from the video signal with the calibration at 704. At 705, a timer is used to prompt re-calibration. For example, the timer may be set to expire in 15 minutes, or one hour, or two hours, or other time durations desired by the caregiver. If the time has expired, the method returns to 701; if not, the method continues to 706, where the video SpO2 value is compared to a threshold to identify changes. If the video SpO2 value crosses the threshold, the method includes sounding an alarm (such as an audible sound and/or a visible alert) at 707, and prompting re-calibration at 701. If not, the method returns to continue measuring at 704. The threshold used to detect a change at 706 can be set by the caregiver to identify changes in video SpO2 that may indicate a clinically significant change in the patient's physiology, for further diagnosis or treatment.

**[0094]** When calibration or re-calibration is not available, the monitor may continue to calculate video SpO2

to identify trends. The trend from the video SpO2 may be used to trigger an alarm when the trend shows that SpO2 is rapidly changing or has crossed an alarm threshold. Clinically relevant patterns (such as repeated desaturations) may also be detected from the video SpO2 signal, between or in the absence of re-calibrations.

**[0095]** When the video-based SpO2 value is calibrated to an accurate measure of oxygen saturation, it can be tracked from there to measure the patient's actual SpO2 value. An example of this is shown in Figure 8, which plots two SpO2 values, one from a traditional contact-based pulse oximeter, and the other from a calibrated video-based pulse oximeter. The video-based SpO2 value in this example is taken from the Red and Green signals, and then calibrated with an absolute SpO2 value as described above. Once calibrated, it is clear from Figure 8 that the video-based SpO2 value tracks the patient's absolute SpO2 value closely. The data presented in Figure 8 was collected during a clinically-relevant desaturation event in which the subject's oxygen saturation dipped and then recovered.

**[0096]** Though the video-based SpO2 measurement can be calibrated from a contact-based pulse oximeter, the video-based SpO2 measurement may exhibit different behavior over time, as compared to a traditional contact-based oximeter. These differences may arise due to the differences in filtering characteristics between the contact-based oximeter and video camera, and/or differences in the light waveforms detected by a remote video as compared to a contact-based sensor, and/or other factors. As an example, the light detected by a remote video camera may be reflected from a shallower depth within the patient's tissue, as compared to contact-based oximetry, which utilizes a contact sensor to emit light directly into the patient's tissue. This difference in the light signal can cause the morphology of the video-detected waveform to differ from a contact-based waveform. As another example, the light detected by a remote video camera is more susceptible to ambient light noise incident on the surface of the region being monitored.

**[0097]** As a result, the SpO2 measurement from the video-detected waveform exhibits some differences from the contact-based SpO2 measurement, even when the two are first calibrated together. An example of this behavior is evident in Figure 8. Between times t1 and t2, the subject's oxygen saturation drops and then recovers to a baseline level BL. Both waveforms track this trend, but the video-based measurement is slower than the contact-based measurement to return to baseline. The result is a difference, labeled  $\Delta S$  (delta saturation) between the two measurements. Because this behavior of the video-based measurement is known, it can be corrected for, by adjusting the value upward during an increasing trend. This adjustment can be tailored based on empirical data. An adjustment may be made by finding the relationship (mapping) between the video-based SpO2 and the contact-based (oximeter) SpO2. This relationship may then be coded within the video system to mimic the oximeter-

based SpO<sub>2</sub>.

**[0098]** In an embodiment, the video-based non-contact monitoring system identifies acute hypoxia in monitored patients, by identifying episodes of decreased oxygen saturation. The system provides continuous monitoring of vital signs such as video-based SpO<sub>2</sub>, rather than discrete, periodic spot-check readings. This continuous monitoring, via either trending or calibrated video SpO<sub>2</sub>, enables the system to identify clinical conditions such as acute hypoxia, and repeated interruptions in air-flow.

**[0099]** In an embodiment, the video-based non-contact monitoring system utilizes a camera that detects light across the visible spectrum. In an embodiment, the camera detects light in only a portion of the visible spectrum, and/or in the infrared spectrum as well.

**[0100]** In an embodiment, a monitoring system is programmed to take certain steps including activating alarms or messages when a suitable physiologic signal is not ascertainable in the field of view. For example, in an embodiment, a processor acquires a physiologic signal (as described above), and determines a physiologic parameter from the signal. However the signal may be lost when the patient moves out of the field of view, or moves in such a way that a physiologic region (such as exposed skin) is not visible, or moves too quickly for accurate tracking. The signal may also be lost if another person or item moves into the field of view and blocks the camera's view of the patient, or if the room becomes too dark (such as if room lights are turned off at night). In any of these or similar situations, the processor starts a timer counting down, and holds the previous value of the calculated physiologic parameter. After a short duration, the processor may send an alert message to be displayed on a screen or otherwise notified to a clinician, to indicate that the signal has been lost and the parameter value is held frozen. If the timer expires, the processor can then sound an alarm or other notification, such as an escalated message or indicator, and remove the frozen physiologic parameter value (or otherwise indicate that it is a previous value, no longer being updated). This can be a system-level alarm or notification, which indicates a problem with the signal acquisition, as distinguished from a physiologic alarm (that would indicate a physiologic parameter of the patient crossing an alarm threshold). This alarm or notification can be a message stating that the room lights have been turned off, or the patient has exited the field of view, or the patient is obscured in the field of view, or the patient is moving, or other applicable circumstance.

**[0101]** This message can be displayed at a remote station (such as a nursing station at a hospital) or on a remote, wireless device (such as a smartphone, tablet, or computer). Additionally, at a central monitoring station (such as a nursing station at a hospital), where display screens display information about multiple different patients, the video-based monitoring system can alert the central station to highlight an individual patient. For ex-

ample, the remote monitoring system can send an alert or flag based on a change in condition (a system-level alarm, a physiologic alarm, an activity level of the patient, etc.), and the central station can then enlarge the video stream from that particular camera. This enables the caregivers at the station to quickly assess the situation in the room and determine if urgent action is needed.

**[0102]** In an embodiment, the processor identifies or is informed that a clinician or caregiver is interacting with the patient, and the processor temporarily halts dynamic tracking of the intensity signal and/or temporarily halts calculation of a physiologic parameter from the intensity signal. This step is taken because such interaction interferes with the camera's view, rendering the light intensity signals more noisy and less reliable. When the interaction is finished, the processor resumes its remote monitoring of the patient.

#### Ambient Light

**[0103]** As mentioned previously, changes in ambient light in the camera's field of view can obscure the subtle variations in the detected pixel streams that are attributable to the patient's physiology. In an embodiment of the invention, a video-based monitoring system includes a calibration strip that can be used to identify and correct for these changes in ambient light. A calibration strip 1100 according to an embodiment is shown in Figure 11. The calibration strip 1100 is sized to fit on the patient (such as along the patient's forehead) and within the field of view 1102 of the camera. In an embodiment, the calibration strip 1100 includes a scale which displays a range of values for measurement, such as a greyscale with two or more grey or white hues; or a color map with two or more different colors. The scale can include a continuous spectrum of varying intensity and/or color, or it can include a set of discrete areas each with a different color and/or intensity. In one embodiment, the color map includes one or more known skin tone colors, which are then compared to exposed skin of the patient to identify an approximation of the patient's skin tone, which can then be used to adjust the exposure settings if the camera based on the light intensity of the skin. These values may vary along one (e.g. longitudinal) or two dimensions of the calibration strip. For example, the calibration strip 1100 shown in Figure 11 includes a grid 1104 with four different discrete regions 1106A, 1106B, 1106C, and 1106D. Each region displays a different intensity and/or color. The colors have a known chromatic value, which allow for the colors in the captured video image to be color balanced to make corrections. Another example is a strip with a grey square or other shape. The intensity of the patch or portions of the patch (such as a grey square) identified in the video image can be used to adjust the exposure settings on the camera. In an embodiment, the calibration strip has a matte finish to reduce reflected light.

**[0104]** In an embodiment, a calibration strip includes

spaces that are Red, Green, Blue, and white. This strip provides a reference for color balancing the region of interest on the patient. For example, if the white space from the calibration strip appears with a green hue on the image, then the region of interest can be color balanced to remove the green skew. This can be particularly helpful for SpO2 measurements.

**[0105]** Figure 9 shows an embodiment of a video-based method of measuring a patient's vital sign. The method includes providing a calibration strip comprising a substrate with a visible scale for viewing by the video camera in the same field of view as the patient, at 901. The method includes detecting, by the video camera, a first light signal reflected from the scale and a second light signal reflected from exposed skin of the patient at 902, and adjusting a calibration of the video camera based on a measurement of the first light signal at 903. The method includes applying the calibration to the second light signal at 904, measuring a vital sign of the patient from the calibrated second light signal at 905, and outputting the measured vital sign at 906 for further processing or display. The scale may be a greyscale or a color map. The measurement of the first light signal can be a measurement of an intensity of the light reflected from the scale, such as a portion of the color map or the greyscale.

**[0106]** The method includes monitoring the measured intensity to detect changes in ambient lighting in the room at 907. For example, at a later time, the system measures a second intensity that differs from the first intensity by a defined amount (such as an amount exceeding a threshold, to avoid excessive adjustments), such as a sudden increase or decrease in intensity due to room lights being turned on or off. If this change meets the defined amount or threshold (or other criteria), the method passes back to 903 to adjust the calibration based on the second intensity, such as by adjusting a coefficient in proportion to the defined amount or the new measured intensity. The coefficient is applied to the light intensity signal, to re-normalize the signal to the second intensity. The coefficient is applied by the camera in generating the red, green, and blue pixel signals. Otherwise, the method returns to 905 to continue monitoring. This method enables the red, green, and blue signals to be normalized to reference values, to better identify the physiologic signals even in differing light conditions. This normalization also enables two different cameras, monitoring two different patients, to be adjusted to the same reference color or brightness values, so that the vital signs or other measurements from the light signals can be compared to each other, without skew due to different camera hardware or light conditions.

**[0107]** In an embodiment, the scale on the patient includes a color map with a plurality of colors, and the measurement of the first light signal is a measurement of a color value of one of the plurality of colors. Then, the calibration is adjusted by comparing the color value to a reference color value and identifying a difference. Base-

line color values from the scale can also be stored at a first time, and then the calibration can be adjusted later based on comparisons of new measurements (of light from the scale) to the stored baseline color values, such as when the new measurement deviates from the baseline by a defined amount. When a second video camera is used in the same room, the second camera can be calibrated based on the same reference, baseline, or other values used for the first camera.

**[0108]** In an embodiment, a calibration strip includes a white space, and the system measures the brightness of that white space in the image. This brightness indicates the amount of light hitting that region. The white space is monitored for changes in brightness, which may indicate a passing shadow or change in lighting conditions, including changes due to movements outside the field of view that change the amount of light reflected onto the region of interest. The color signals from the region of interest can then be filtered according to the changes in brightness, to continue tracking SpO2 (or another vital sign) during transient changes in lighting, such as due to motion in the room. This can be done, for example, with an adaptive filter based on the reference signal from measurement of the white space. Average light intensity within the identified white space can be used as a baseline to compensate non-physiological changes in the sampling regions. Alternatively, the color signals and/or vital sign measurements can simply be discarded during these transient changes.

**[0109]** In an embodiment, a calibration strip includes a graphic with high contrast, such as a dark dot, cross or circle on a white space, or a grid of colors. The system can track this high contrast shape to track movement of the patient. For example, the system can track the position and orientation of the high contrast graphic, and can generate a motion signal that tracks that movement. The motion signal may be a transform that maps the movement of the graphic. The same transform is applied to the region of interest in the patient, to track the movement of that region. If the transform reveals that the region of interest has exited the field of view, then a new region of interest is identified, based on the desired vital sign. Further, limits can be placed on the allowable rate of motion (such as angular rotation limits), and if the limits are exceeded, the color signals and/or measurements from the region of interest can be discarded.

**[0110]** In another embodiment, the calibration strip includes light emitters that emit light of selected wavelengths into the patient, but without the detector of traditional contact-based oximetry sensors, and without any transmitter for transmitting detected light. For example, a calibration strip 130 according to another embodiment is shown in Figure 12. In this embodiment, the calibration strip 130 includes an adhesive patch 132 that is sized to fit within the field of view of a non-contact camera, such as on the patient's skin. The patch 132 includes a top surface 134 that faces away from the patient, opposite a bottom surface 136 that faces toward and contacts the

patient. The top surface 134 carries a scale or graphic 138. The bottom surface carries an adhesive 140 that removably adheres to a patient's skin. The patch 132 also includes two emitters 142, 144 coupled to a battery 146 and a microprocessor 148. When the patch is placed on a patient's skin, the processor 148 drives the emitters 142, 144 to emit light sequentially into the patient's skin. The processor drives the emitters in a four-part sequence, in which the first emitter is on, then both emitters are dark, then the second emitter is on, and then both emitters are dark. This sequence is repeated at high frequency, such as 15-30 Hz.

**[0111]** However, notably, the patch 132 does not include a photodetector or any type of transmitter. Rather, the detector is a non-contact video camera viewing the patient, as described above. The video camera records image frames that include at least the portion of the patient's skin surrounding or near the patch 132. Light from the emitters travels through the patient's tissue and out through this portion of the patient's skin, such that it can be detected by the video camera. This system is a hybrid approach, employing contact-based emitters and a non-contact, remote detector. The system benefits from having dedicated light emitters at chosen wavelengths (for example, a narrow range of red and green wavelengths), creating a stronger physiologic signal in the detected image frames, while at the same time avoiding the drawbacks of a tethered sensor system. The patch 132 does not have any cables or wires connecting it to a monitor, nor any wireless communication. The patch 132 does not require any communication at all between itself and the camera or the monitor (such as the camera 214 and monitor 224 in Figure 2A). As a result, the patch can omit components such as a wireless transmitter or receiver and supporting components such as batteries for those devices. The processor 148 carried by the patch can operate at very low power, operating only to drive the emitters 142, 144 and not to process or transmit any detected signal. The processor and emitters can be powered by a small battery 146. The patch is also small and lightweight, making it relatively comfortable for the patient to wear, and it does not interfere with the patient's mobility. The camera may begin monitoring the patient's vital signs automatically when it detects the emitted light, or it may be turned on by a caregiver.

**[0112]** It should be noted that the scale 138 shown in Figure 12 is optional. In an embodiment, the patch 132 omits the scale 138 on the top surface, and is not used as a calibration strip. In another embodiment, the patch 132 includes a single color on the top surface, such as white, for use in measuring brightness and detecting passing shadows.

#### Independent Component Analysis

**[0113]** Due to the exposure of the camera detector to significant ambient light noise, the video-based system employs new approaches to filter the ambient light noise

and identify the physiologic signal from which the patient's vital sign can be measured. An approach for filtering according to an embodiment of the invention is demonstrated in Figures 13-16. In this embodiment, independent component analysis (ICA) is used to decompose the Red, Green, and Blue pixel streams into individual components. ICA is a filtering method that, based on certain assumptions, separates input signals (also called source signals) into separate, independent components that are mixed together in the input signals. The ICA method is described in detail in the following paper: Hyvärinen, A., & Oja, E. (2000). Independent component analysis: algorithms and applications. *Neural networks*, 13(4), 411-430.

**[0114]** In the context of video-based monitoring, the source signals are the Red, Green, and Blue pixel streams, and the independent components are the heart rate and the noise. Referring to Figure 13, the source signals are shown in the three plots on the left, with the Red pixel stream on top, Green in the middle, and Blue on the bottom. These source signals are decomposed via an ICA method into three separate, independent components, shown on the three plots on the right (labeled Component 1, Component 2, and Component 3).

**[0115]** As shown in Figure 13, Component 1 exhibits a repeating pattern of modulations at a relatively steady frequency. Component 1 is constructed from the portions of the source signals that modulate at that frequency. In this case, the frequency of the modulations in Component 1 represents the heart rate of the patient. The contributions of the patient's heart rate to each source signal have been pulled together and combined into the waveform of Component 1, creating a waveform that identifies the heart rate more clearly than any single source signal did. The patient's heart rate can be measured from the primary frequency of Component 1.

**[0116]** Still referring to Figure 13, Components 2 and 3 are relatively more erratic, and do not exhibit a clear primary frequency. These components capture the noise that corrupted the Red, Green, and Blue source signals. Each of the source signals represents a different mixture or combination of Components 1, 2, and 3.

**[0117]** By utilizing ICA to decompose the source signals, an underlying physiologic signal such as heart rate or respiration rate can be identified. As discussed above, different groups of pixels or regions can be selected to measure different vital signs, such as heart rate and respiration rate. Figure 13 represents the source signals from a first group of pixels or regions that modulate with the patient's heart rate. These signals are decomposed via ICA to arrive at a relatively clean heart rate signal in Component 1. A different group of pixels or regions that modulate with respiration rate can also be decomposed via ICA to arrive at a relatively clean respiration rate signal. Another region may be decomposed via ICA to arrive at a pulsatile signal that demonstrate perfusion status of the patient (such as by Delta POP or DPOP, by measuring the variations in amplitude of the pulses at the top

and bottom of the baseline modulations). These vital signs may be measured from the same region or different regions.

**[0118]** In Figure 13, Component 1 exhibits the most regular frequency, as shown by the plotted vertical lines. Vertical lines are placed in the plots of Components 1, 2, and 3 at each local maximum in the waveforms. The component with the most regularly spaced vertical lines is chosen as the component that represents the patient's heart rate. In Figure 13, this is clearly Component 1.

**[0119]** Figure 14 shows the results of an ICA method applied to a video stream to measure a patient's heart rate. The figure shows heart rate calculated by a traditional contact-based oximeter (solid line) as well as heart rate from an ICA filtered video stream (x's) from the same subject over the same time duration. The ICA-based heart rate shows good correlation with the traditional oximeter values.

**[0120]** After decomposing the source signals via ICA to identify a physiologic component (such as heart rate), that component can then be used to filter the original input signals, as shown in Figure 15. Figure 15 shows three traces - Component 1 on top, the Red source signal in the middle, and the Green source signal on the bottom. The vertical lines marked on the local maximums of Component 1 are projected onto the Red and Green pixel streams. The locations of these projections signify heart beats in the Red and Green pixel streams, even when these source signals are corrupted by noise. The ICA-derived heart rate signal of Component 1 can be used to identify the location of individual pulses in the source signals. The ICA technique finds the best representative pulse signal, which can then be used to locate pulses in the original Red and Green source signals.

**[0121]** Figure 16 depicts a flowchart of a method for measuring a vital sign of a patient with ICA, according to an embodiment. The method includes receiving a video signal from a video camera at 1601, and extracting from the video signal the source signals at 1602, such as time-varying red, green, and blue signals. The method includes performing ICA to generate a component signal having a primary frequency at the heart rate of the patient at 1603. Performing ICA involves decomposing at least two of the source signals into component signals, and selecting the component signal that exhibits the contribution of the patient's heart rate, as explained above. The method then includes identifying, in the selected component signal, an individual pulse representative of an individual heart beat at 1604. The method includes locating a corresponding portion of at least two of the red, green, and blue source signals at 1605. This can be done by determining a fiducial in the component signal (e.g. the maxima, minima, peak of the first derivative, etc.), and using this fiducial to identify a corresponding pulse or location in the source signals. Then, for each of those two source signals, the method includes adding the located portion to a weighted average pulse at 1606. This produces at least two of a red weighted average

pulse, a blue weighted average pulse, and a green weighted average pulse. The method then includes measuring blood oxygen saturation of the patient from the weighted average pulses of the two signals at 1607 (such as by measuring an ROR, and computing SpO2 from the ROR). Heart rate can also be measured from the primary frequency of the component signal. The vital signs are output for further processing or display. In an embodiment, the method also includes triggering an audio beep at 1608 in synchrony with the individual pulse identified in the component signal, or in synchrony with the located corresponding portion of one or two of the color signals. This audio beep signifies the occurrence of a cardiac pulse. Instead of an audio beep, other audible or visual alerts may be triggered or displayed.

**[0122]** The ICA-derived pulsatile component signal is thus used as a trigger to inform the processor where to look in the original signals for relevant physiologic information. In turn, this trigger can be used to control an ensemble averaging method, in which sequential pulses are averaged with past pulses to create a smoother average cardiac pulse for each source signal. The ICA-derived trigger may also be passed to another medical device, such as a pulse oximeter, blood pressure monitor, or other monitor or processor, to inform that device that a cardiac pulse has been detected and the time or location of that pulse.

#### Noise Reduction

**[0123]** Another way to address noise is to identify non-physiologic peaks within the frequency domain, and remove those from the video signals. Two methods for identifying non-physiologic peaks are summarized here.

**[0124]** In one method, in the frequency domain, peaks are identified that remain stationary over a duration of time. Over a sufficient period of time (long enough for a few cycles of the vital sign - for example, 5-10 seconds for heart rate, or 20-30 seconds for respiration rate), peaks that remain stationary are likely to be non-physiological, such as peaks caused by aliasing from flickering room lights, while physiologic peaks should move and shift with the patient's state. A frequency transform such as an FFT can be performed over different time durations (such as different window sizes), and the frequencies that remain stationary, by appearing the same regardless of window size, are likely to be non-physiological. These identified frequencies can be removed by filtering. A flowchart illustrating this method is shown in Figure 17A. In an embodiment, the method includes performing a frequency transform over first and second time windows of different sizes (different time durations) at 1701. The method includes comparing frequency peaks in the transforms at 1702, and identifying stationary frequency peaks at 1703. The method then includes filtering the video (source) signal(s) to remove the stationary frequency at 1704.

**[0125]** The number of window sizes, and their sizes

relative to each other, can be varied to achieve a desired result. In an embodiment, two different window sizes are used, one 20 seconds in duration and the other 10 seconds in duration. In another embodiment, three window sizes are used, such as 20, 10, and 7 seconds each. This analysis can be done on each pixel signal individually, to remove identified frequencies from each signal, or it can be done on one signal and then the identified frequencies can be removed from all signals.

**[0126]** In another method, in the frequency domain, peaks are identified that move based on frame rate. Frequency peaks that move to another position or disappear when the video frame rate is adjusted may be taken as non-physiological, because physiologic modulations do not disappear or move instantaneously based on the video characteristics. In an embodiment, the frame rate sweeps at a constant sweep rate over a range of frequencies, or moves along a trajectory (such as a first frame rate for a first time duration, then a second frame rate for a second time duration, etc), and frequency peaks that move with that sweep or trajectory are considered non-physiological. Frequency peaks that move at the sweep rate are particularly suspect and can be removed. The speed of the sweep is faster than the expected variation of physiological parameters, such as heart rate. The frame rate can also change in random or pseudo-random ways, or through a set of non-stationary values, such as three or more discrete, different frame rates. Further, a frequency peak that remains stationary upon the change in frame rate is more likely to be physiological. A stationary peak can be identified, and a vital sign such as heart rate measured from this stationary peak. A flow-chart illustrating this method is shown in Figure 17B. In an embodiment, the method includes adjusting the frame rate of the video signal at 1711, and identifying peaks in the frequency domain that change or move with the adjusted frame rate at 1712. The method then includes filtering the source signal(s) to remove the identified frequency at 1713. In an embodiment, after the noise frequency has been identified, the frame rate can be fixed, until a later time when it is varied again to re-check for noise peaks.

**[0127]** The particular range of frame rates may depend on the capabilities of the camera hardware, and the light conditions. In an embodiment, the frame rate is varied from the highest well-exposed frame rate to lower frame rate, in one or more steps. An example range is 10-25 frames per second. In an embodiment, the period of time during which the frame rate is varied is longer than the expected period of the physiologic frequency (such as heart rate). The analysis described in Figures 17A and 17B can be done on each pixel signal individually, to remove identified frequencies from each signal, or it can be done on one signal and then the identified frequencies can be removed from all pixel signals.

## Optical Splitter

**[0128]** In another embodiment, an optical splitter is employed in order to obtain two light signals from a single camera. These two light signals encompass the same field of view, monitoring the same subject, over the same time period, but the two signals can be filtered differently to facilitate physiological measurements. The two signals are synchronized in time and field of view, and include the same noise components, so the same de-noising operations can be used on both signals. The optical splitter is a simpler solution than two separate cameras, and provides more information than a single camera.

**[0129]** An optical splitter 1810 according to an embodiment is shown in Figure 18. The optical splitter 1810 is used to split a single light signal into two signals that pass through two different filters. The filters are chosen based on the physiologic signal that is to be measured. For example, for SpO<sub>2</sub>, the filters are chosen based on the extinction coefficients of hemoglobin. The two filters can pass visible and non-visible light, respectively, such as red and infrared light, or two narrow ranges of visible light. One filter may pass a narrow range of red wavelengths, and the second filter may pass a narrow range of green wavelengths, to mimic the red and infrared signals emitted by traditional contact pulse oximetry emitters. Referring to Figure 18, the optical splitter includes an aperture 1812 that receives an incoming light signal 1800. The optical splitter includes a beam splitter 1814 positioned behind the aperture, in the path of the incoming light. The beam splitter 1814 divides the incoming light signal 1800 into two signals 1800A and 1800B. An example of a beam splitter is a dielectric mirror or a beam splitter cube, operating to split the incident light into two or more paths, in not necessarily equal proportions of strength. The separated light signal 1800B passes to a mirror 1816 that re-directs the light signal into the camera. Each light signal passes through a respective filter 1820A, 1820B. In the example of pulse oximetry, the filter 1820A is designed to pass a narrow range of red wavelengths, while the filter 1820B passes a narrow range of green wavelengths. The filtered light signals are received by respective detectors or light sensors 1822A, 1822B that register the resulting images. The result is two time-varying image signals filtered for specific wavelengths. Regions of interest are identified in the two signals for the calculation of vital signs such as SpO<sub>2</sub>, heart rate, and respiration rate, as described above. Absolute SpO<sub>2</sub> can be calculated via a pre-calibrated look-up table, without the need for periodic re-calibration via oximeter spot check.

**[0130]** In another embodiment, additional splitters may be used to divide the light into more than two beams, to pass through additional filters chosen for other physiologic parameters. For example, an additional beam 1800N can be passed through a filter chosen for the measurement of total hemoglobin. In another example, a filter is chosen for carboxyhemoglobin, or for methemoglobin.

moglobin. In an embodiment, the filters are arranged on a rotating wheel, so that they are rotated in and out of the path of the light 1800, 1800A, or 1800B, to filter the incoming light as needed for the measurement of the physiologic parameters. This mechanical filter actuator can select appropriate filters to measure different parameters from the patient at different times.

#### Data

**[0131]** Non-contact video-based monitoring methods have been employed in various test environments to confirm their utility. Some of that testing is summarized below. For example, Figures 19, 20, and 21 show video-based measurements of heart rate, respiration rate, and SpO<sub>2</sub> as compared to reference measurements, during a clinical study. During the study, the reference measurements were taken as follows: for heart rate and SpO<sub>2</sub>, from a contact-based pulse oximeter, and for respiration rate, from a ventilator. A video camera was spaced apart from and oriented at the animal subject, and video signals were captured through the course of an oxygen desaturation. The video signals were used to calculate heart rate, respiration rate, and SpO<sub>2</sub>, and these measurements were compared to the reference measurements, as shown in Figures 19, 20, and 21. These figures show good agreement between the video-based measurements and the reference measurements. Two separate regions of interest on the skin were chosen, one for the determination of rates (RRvid and HRvid) and the other for the determination of saturation (SvidO<sub>2</sub>).

**[0132]** Figure 19 is a scatter plot of the video-based heart rate measurements on the y-axis, against the reference heart rate measurements on the x-axis (both in beats per minute). The dotted line is a least squares fitted regression line. The expected 1:1 correspondence line is also shown, but is mostly hidden by the regression line, showing very good fit between the two. Each desaturation episode is shaded separately.

**[0133]** Figure 20 is a scatter plot of the video-based respiration rate measurements against the reference respiratory rate measurements from the ventilator (both in breaths per minute). The dotted line is a least squares fitted regression line. The expected 1:1 correspondence line is shown in solid black. The size of each circle corresponds to the number of data points at that location; this visualization was required due to many co-located data points in the plot.

**[0134]** Figure 21 is a scatter plot of the video-based SpO<sub>2</sub> measurements against the reference SpO<sub>2</sub> measurements (both in %). The dotted line is a least squares fitted regression line. The expected 1:1 correspondence line is shown in solid black. Each desaturation episode is shaded separately. Changes in oxygen saturation were calculated using a ratio of ratios derived from the red (R) and green (G) signals, where the two signals were first normalised by dividing their cardiac pulse amplitudes by the signal baseline values. As discussed above, using a

standard RGB camera only allows for a relative saturation value to be determined from this normalised ratio of the amplitude of two of the signals. Hence this required calibration against known values from the reference pulse oximeter to provide an absolute value of SvidO<sub>2</sub>.

**[0135]** The systems and methods described here may be provided in the form of tangible and non-transitory machine-readable medium or media (such as a hard disk drive, hardware memory, etc.) having instructions recorded thereon for execution by a processor or computer. The set of instructions may include various commands that instruct the computer or processor to perform specific operations such as the methods and processes of the various embodiments described here. The set of instructions may be in the form of a software program or application. The computer storage media may include volatile and non-volatile media, and removable and non-removable media, for storage of information such as computer-readable instructions, data structures, program modules or other data. The computer storage media may include, but are not limited to, RAM, ROM, EPROM, EEPROM, flash memory or other solid state memory technology, CD-ROM, DVD, or other optical storage, magnetic disk storage, or any other hardware medium which may be used to store desired information and that may be accessed by components of the system. Components of the system may communicate with each other via wired or wireless communication. The components may be separate from each other, or various combinations of components may be integrated together into a medical monitor or processor, or contained within a workstation with standard computer hardware (for example, processors, circuitry, logic circuits, memory, and the like). The system may include processing devices such as microprocessors, microcontrollers, integrated circuits, control units, storage media, and other hardware.

**[0136]** Although the present invention has been described and illustrated in respect to exemplary embodiments, it is to be understood that it is not to be so limited, since changes and modifications may be made therein which are within the full intended scope of this invention as hereinafter claimed.

#### Claims

1. A video-based method of measuring a patient's vital sign, comprising:

receiving a video signal from a video camera, the video signal having a field of view that includes exposed skin of a patient;  
extracting from the video signal a time-varying color signal for each of a plurality of regions, each region including exposed skin of the patient;  
identifying a frequency content of each color signal;

- selecting two or more non-adjacent regions that have a shared frequency content comprising a modulation at a shared frequency;  
combining the color signals of the selected regions;  
measuring a vital sign from the combined color signal; and  
outputting the vital sign for further processing or display.
- 2. The method of claim 1, wherein combining comprises averaging the color signals from the selected regions.
- 3. The method of claim 2, wherein the selected regions have different sizes.
- 4. The method of any preceding claim, wherein the vital sign comprises heart rate, and optionally:  
  
measuring the heart rate comprises accumulating frequency peaks from the combined signal, selecting a median frequency, and updating a running average heart rate from the selected frequency.
- 5. The method of claim 3, further comprising updating the selected regions by adding or removing regions from the selected regions based on the frequency content of each region.
- 6. The method of any preceding claim, further comprising calculating a statistic of the combined color signal, wherein the statistic comprises an amplitude, a variability, a skew, or a signal to noise ratio.
- 7. The method of claim 6, wherein at least one of a) and b):  
  
a) combining the color signals comprises applying a weight to each color signal that is being combined, the weight being based on the statistic;  
b) the method further comprises determining a weight for the vital sign, the weight being based on the statistic, and adding the vital sign to a running average based on the weight.
- 8. The method of claim 2 or any of claims 3 to 7 that depend on claim 2, wherein selecting the two or more non-adjacent regions further comprises selecting regions that satisfy a quality criterion, and wherein the quality criterion comprises a signal to noise ratio that satisfies a threshold, or a skew value that satisfies a threshold, and optionally:

the method further comprises removing a color

signal from the combined color signal upon failure of the removed color signal to satisfy the criterion.

- 5 9. The method of any preceding claim, wherein extracting the time-varying color signal comprises extracting two time-varying color signals for each region, and wherein measuring the vital sign comprises measuring oxygen saturation from the two time-varying color signals.
- 10 10. A camera-based method of measuring a patient's vital sign, comprising:  
  
receiving an image signal from a camera having a field of view encompassing exposed skin of a patient, the image signal comprising a sequence of image frames over time;  
within the field of view, identifying a plurality of regions;  
for each region of the plurality, extracting an intensity signal comprising a time-varying light intensity detected in the region;  
selecting regions whose intensity signals comprise a modulation at a shared frequency, wherein the selected regions are non-adjacent in the image frame;  
combining the intensity signals of the selected regions to produce a combined intensity signal; and  
measuring heart rate of the patient from the combined intensity signal.
- 11 11. The method of claim 10, wherein the measured vital sign is a first measured vital sign, and wherein the method further comprises:  
  
selecting a second set of regions whose intensity signals comprise a second modulation at a second, different, shared frequency, wherein the second set of selected regions are non-adjacent in the image frame;  
combining the intensity signals of the second set of selected regions to produce a second combined intensity signal; and  
measuring respiration rate of the patient from the second combined intensity signal, further wherein optionally:  
  
the method comprises displaying the image signal on a display, highlighting on the display the selected regions for the first vital sign, and highlighting on the display the second, different set of selected regions for the second vital sign.

12. The method of claim 10 or 11, wherein at least one

of c) and d):

c) the modulation at the shared frequency is above an amplitude threshold;

d) extracting the intensity signal for each region comprises extracting a red intensity signal comprising a time-varying average of red values in the region, a green intensity signal comprising a time-varying average of green values in the region, and a blue intensity signal comprising a time-varying average of blue values in the region.

f) displaying a number derived from the first measured vital sign, and triggering an alarm based on the displayed number and stored alarm conditions, wherein the stored alarm conditions omit a sensor-off or sensor-disconnect alarm.

13. A camera-based method of measuring a patient's vital sign, comprising:

receiving an image signal from a camera having a field of view encompassing exposed skin of a patient;

receiving a user input identifying a region in the image signal;

measuring a first vital sign of the patient, comprising:

from the region, selecting a first set of pixels that exhibit a first modulation at a first shared frequency, wherein the first selected pixels are non-contiguous; and

measuring the first vital sign of the patient from the selected first pixels; and outputting the first measured vital sign for further processing or display.

14. The method of claim 13, further comprising:

measuring and outputting a second vital sign of the patient, comprising:


from the region, selecting a second set of pixels that exhibit a second modulation at a second shared frequency different from the first shared frequency, wherein the second selected pixels are non-contiguous;

measuring the second vital sign of the patient from the selected second pixels, wherein optionally:

the first vital sign comprises heart rate and the second vital sign comprises respiration rate.

15. The method of claim 13 or 14, further comprising at least one of e) and f):

e) receiving a second user input moving the region to a new location, and subsequently measuring the first vital sign from the new location;



**FIG.1**  
PRIOR ART




FIG. 2A




FIG. 2B




FIG.3A




FIG.3B




FIG.4B




FIG.4A




FIG. 5A




FIG. 5B




FIG. 5C




FIG. 5D




FIG.5E

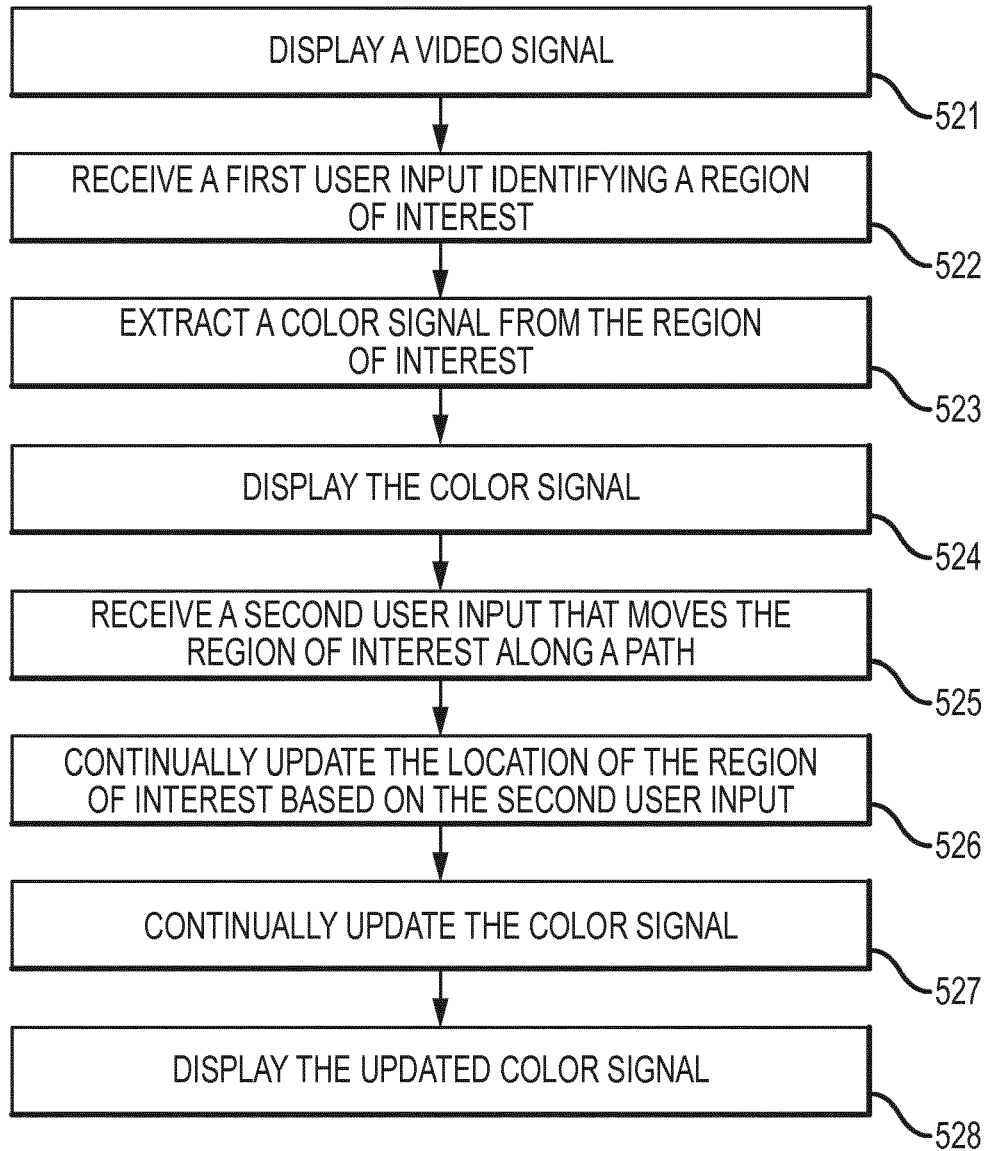


FIG.5F

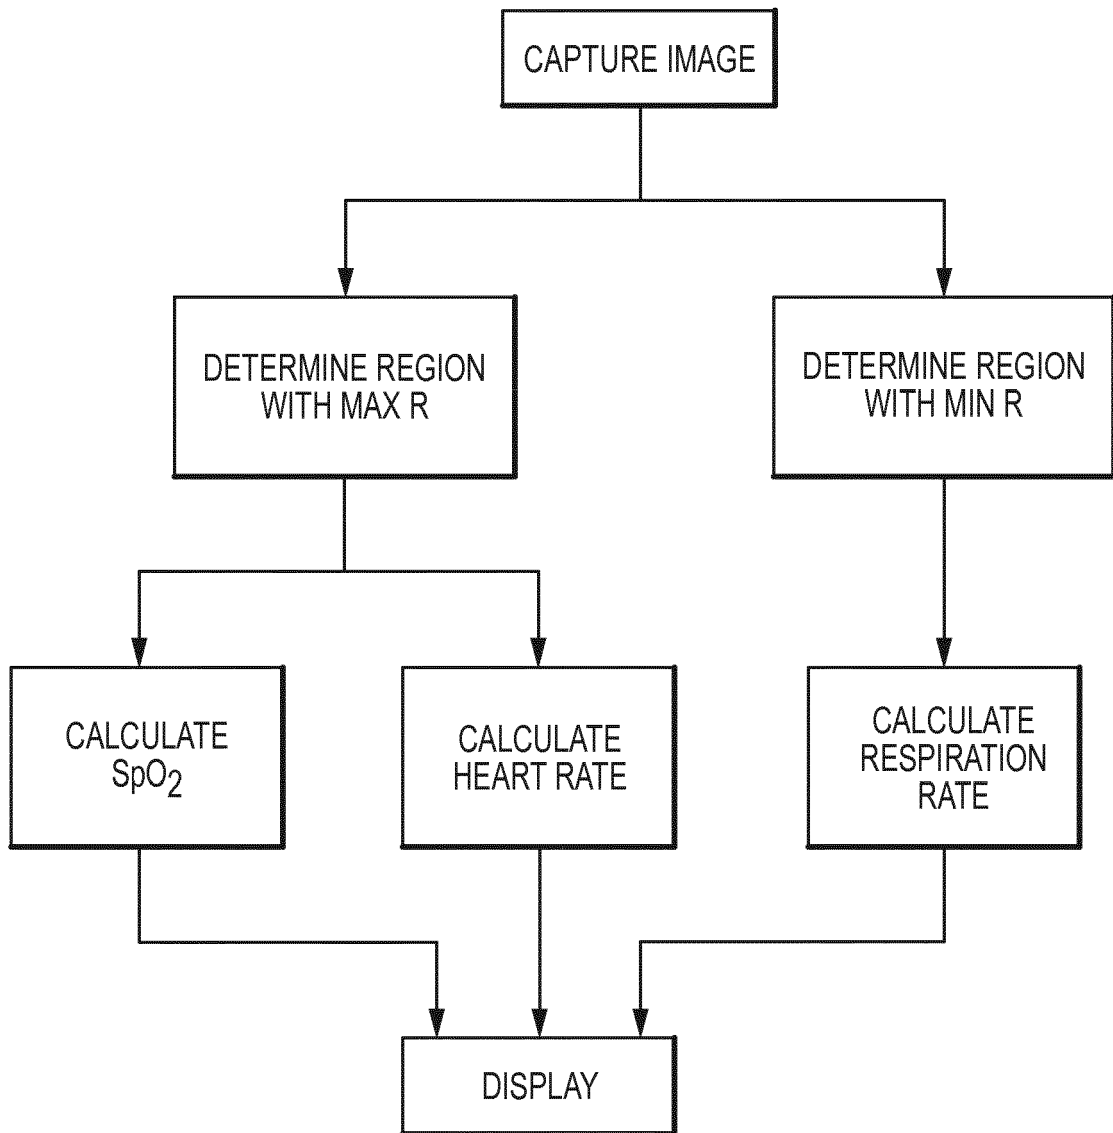


FIG.6A




FIG.6B




FIG. 7




FIG.8




FIG.18

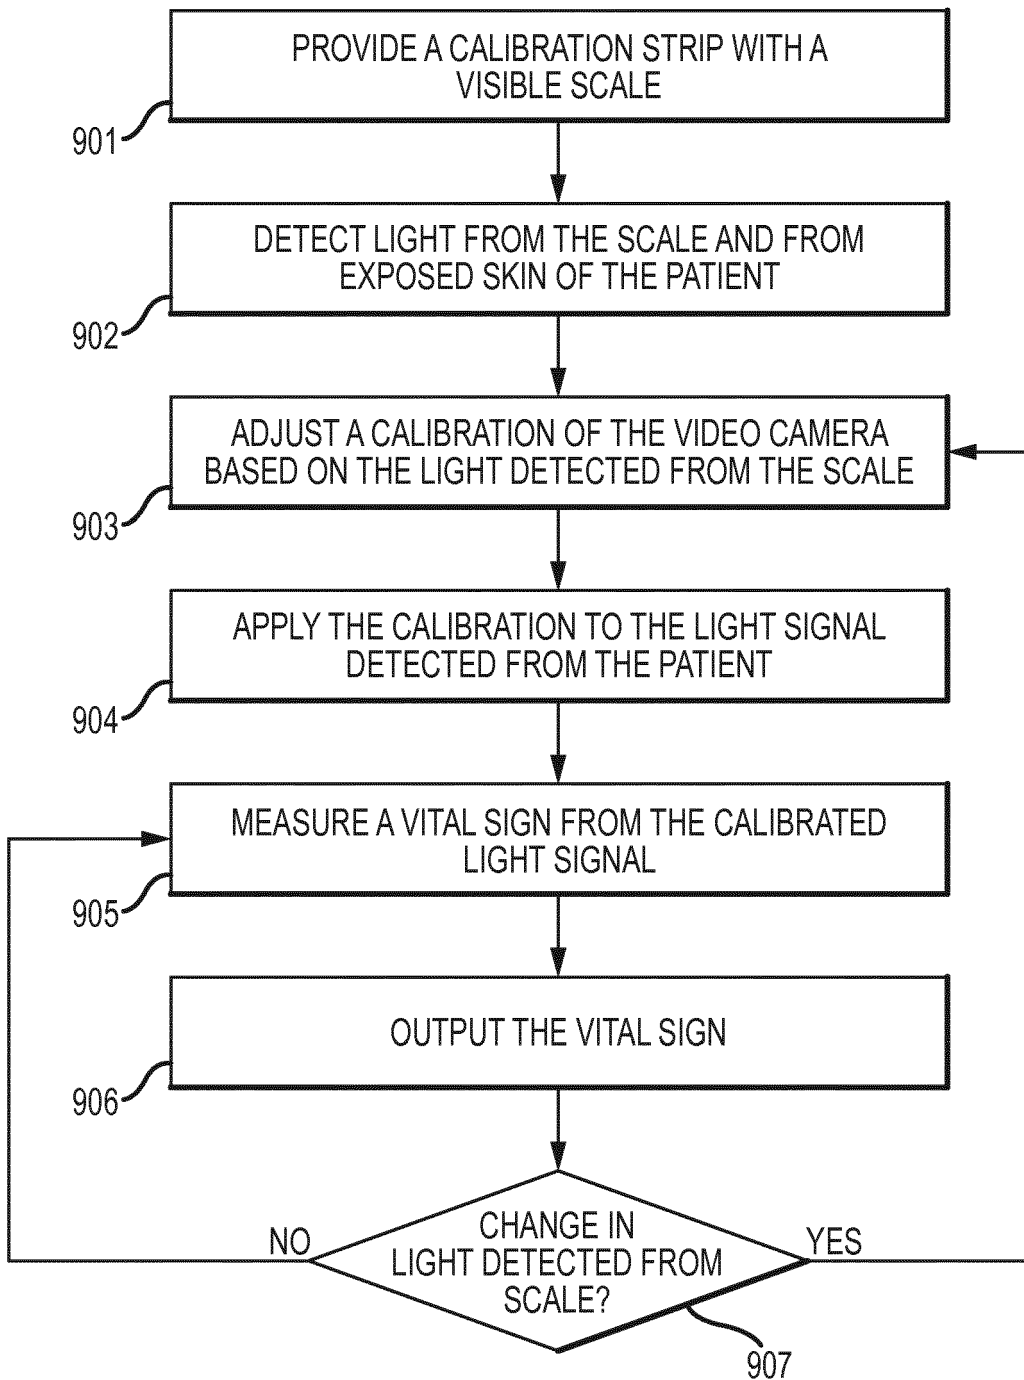


FIG.9




FIG.10A

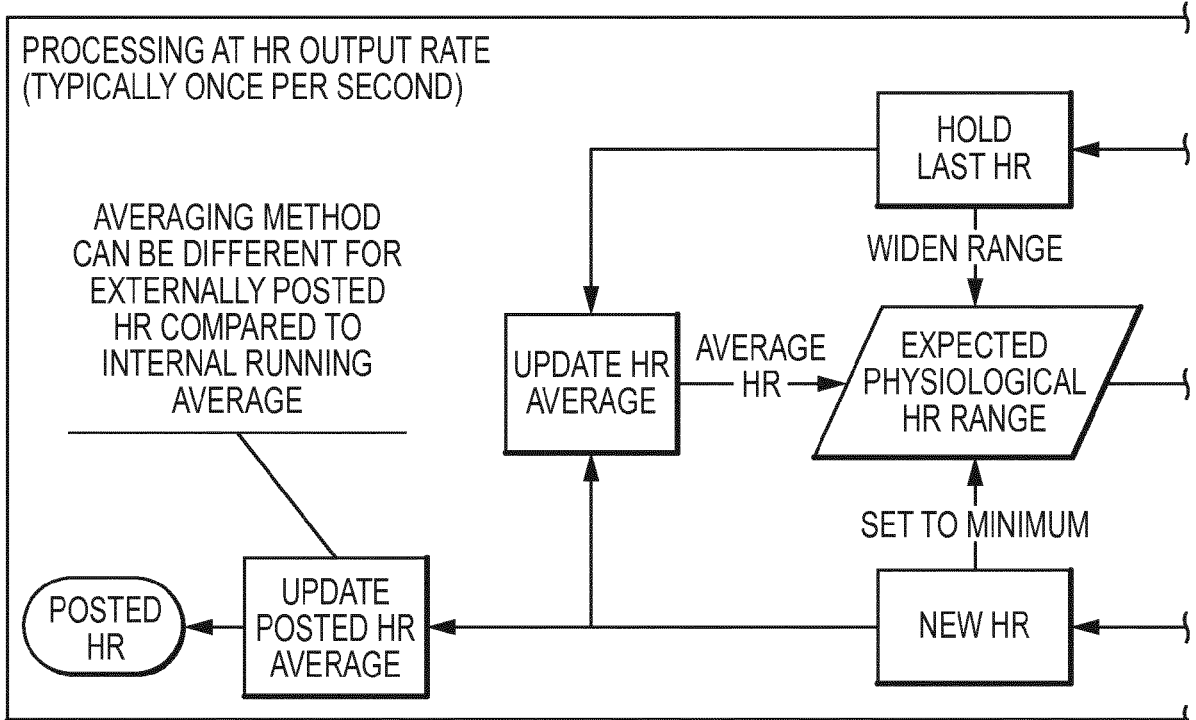
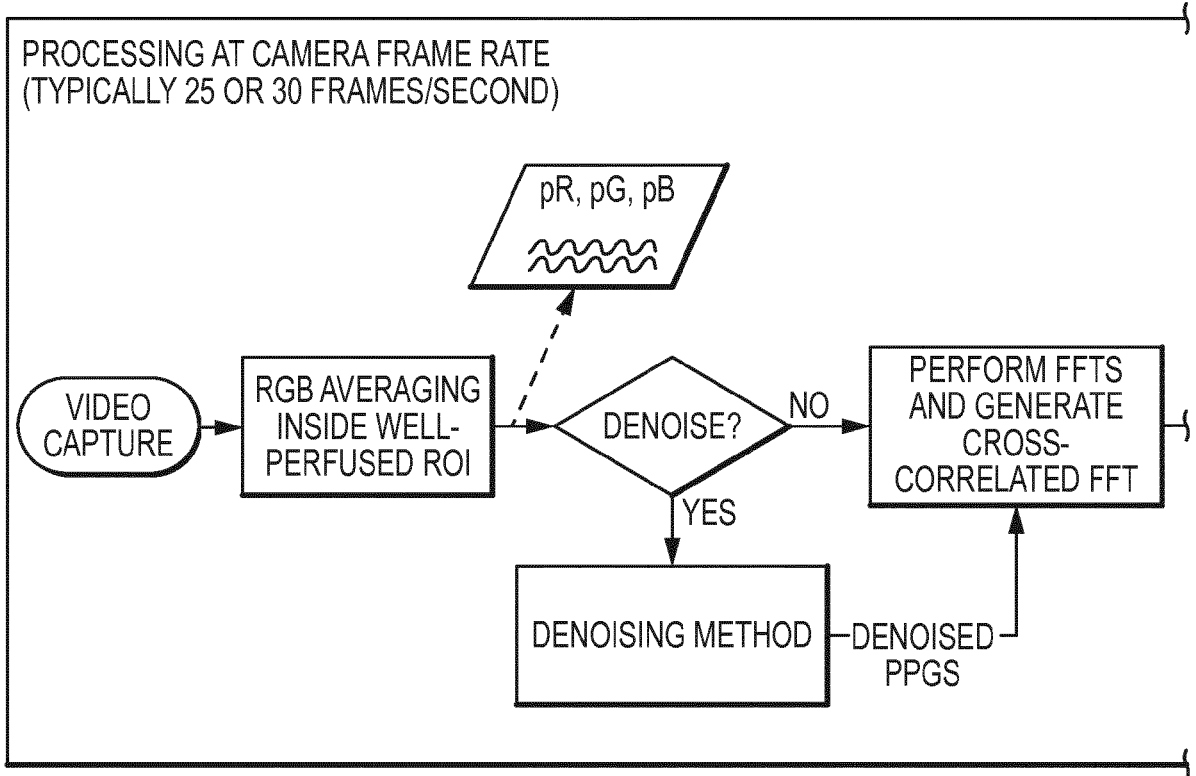


FIG.10B

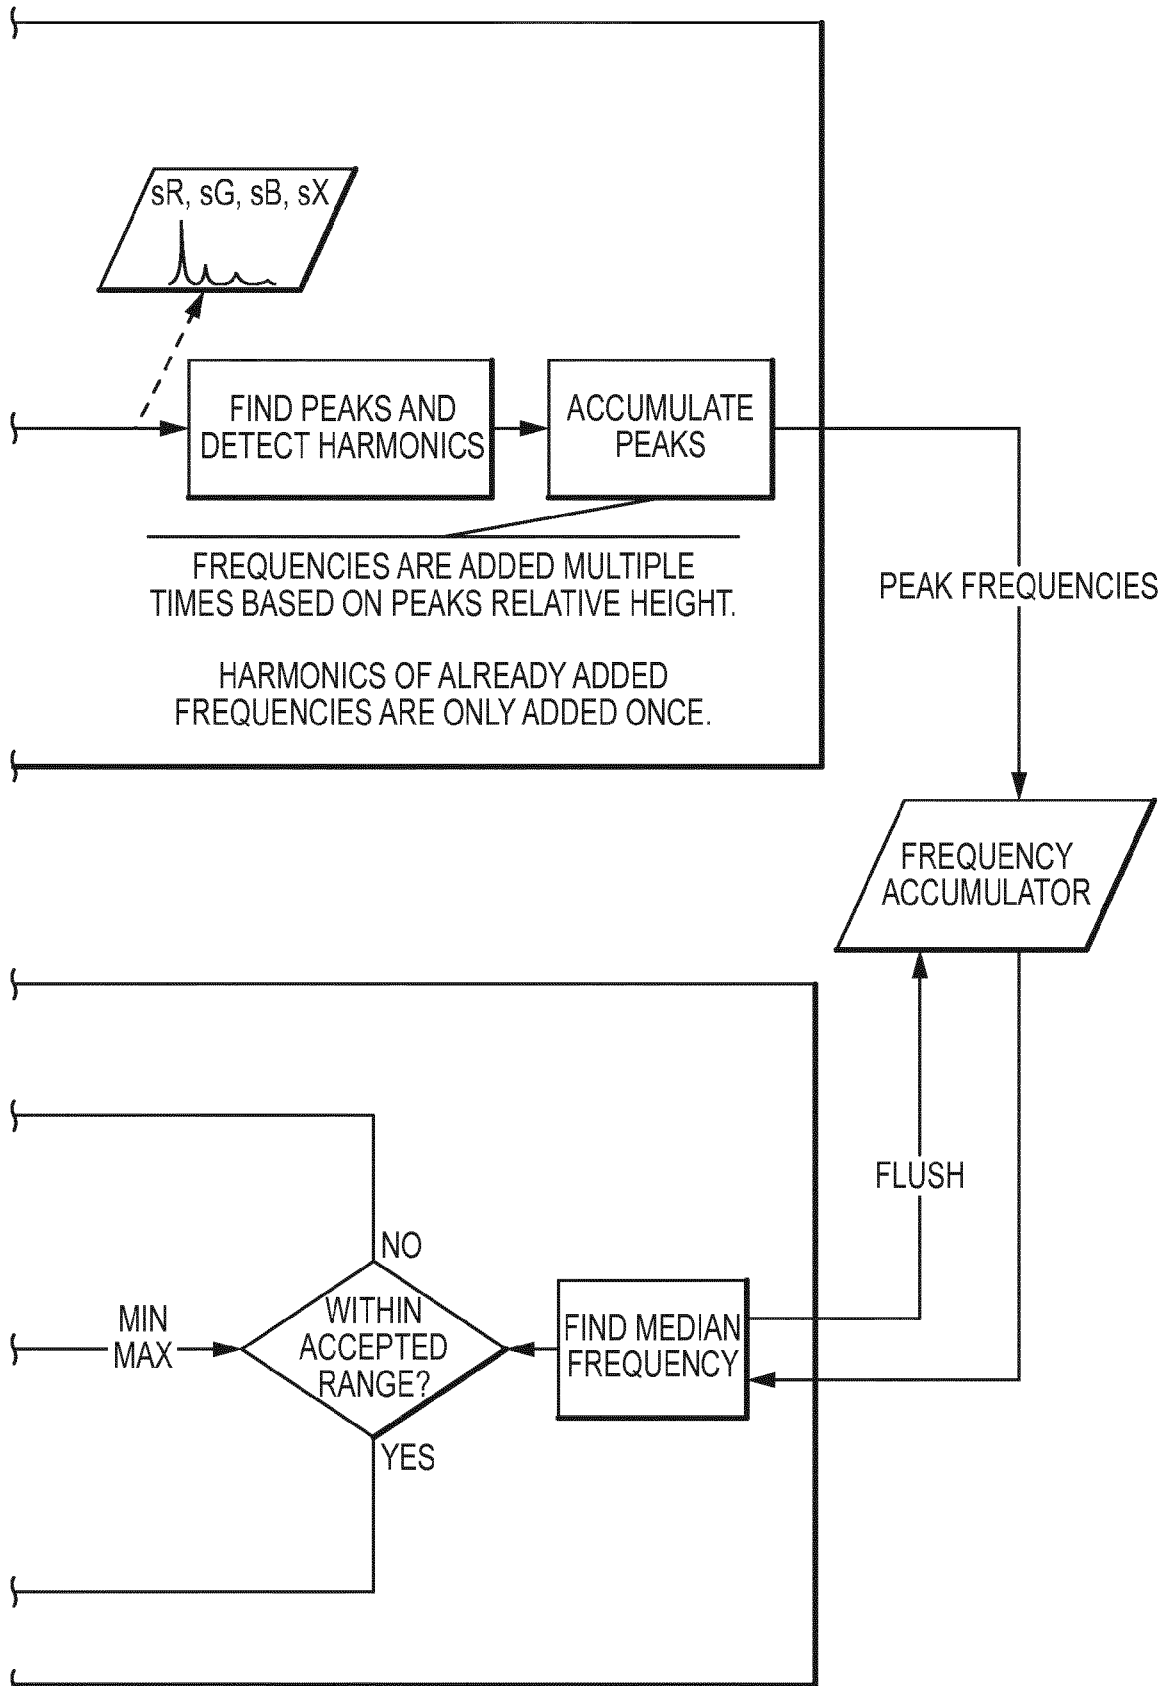


FIG.10B-1




FIG. 11




FIG. 12

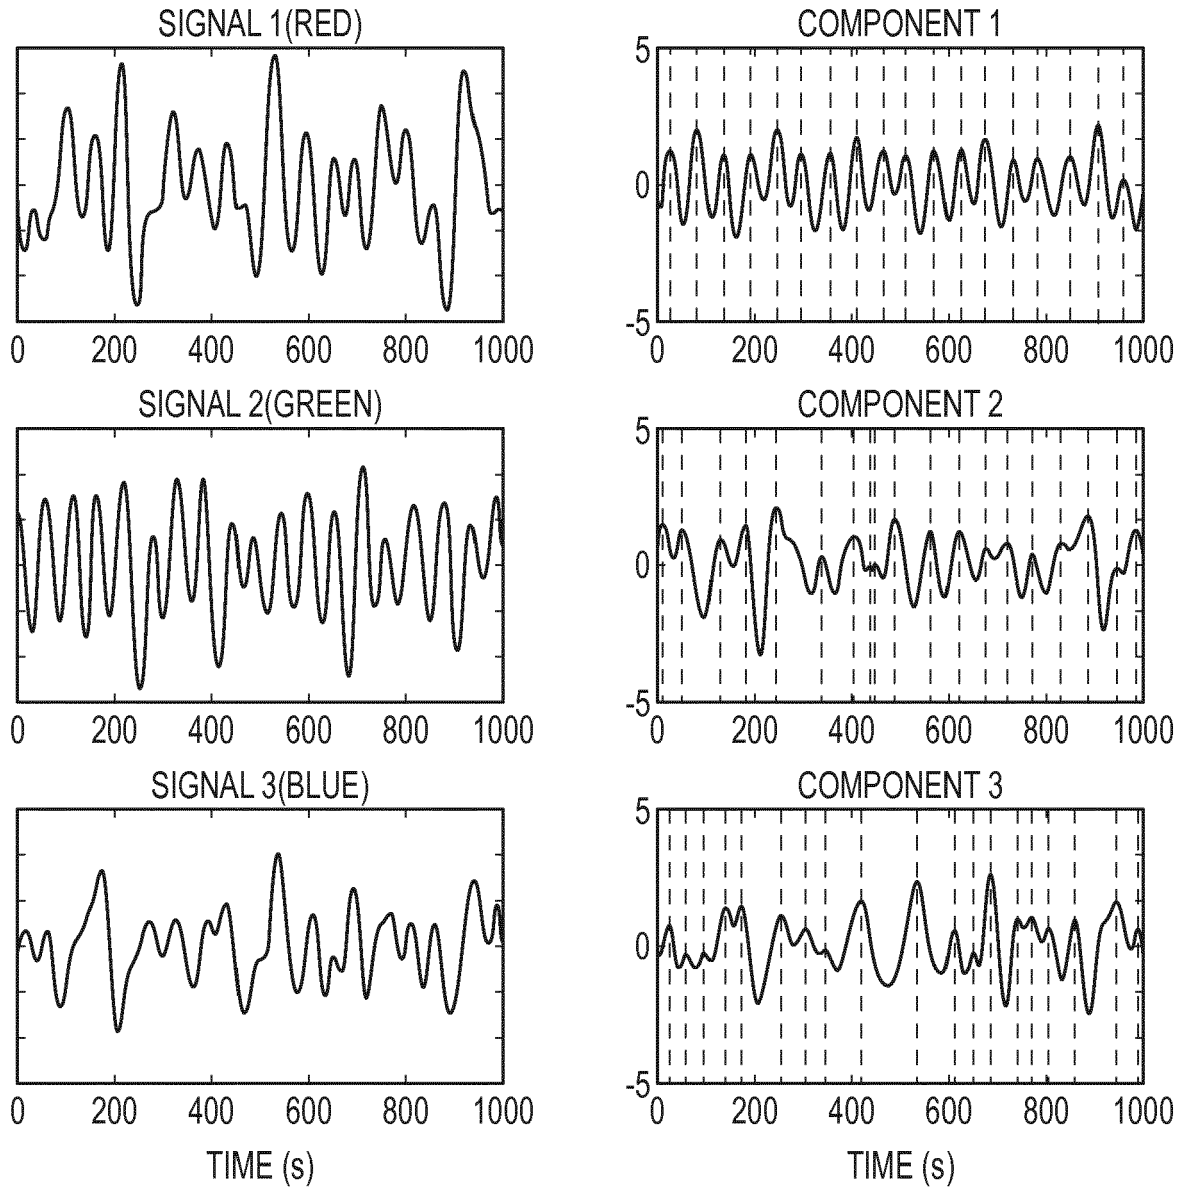


FIG.13




FIG. 14




FIG. 15

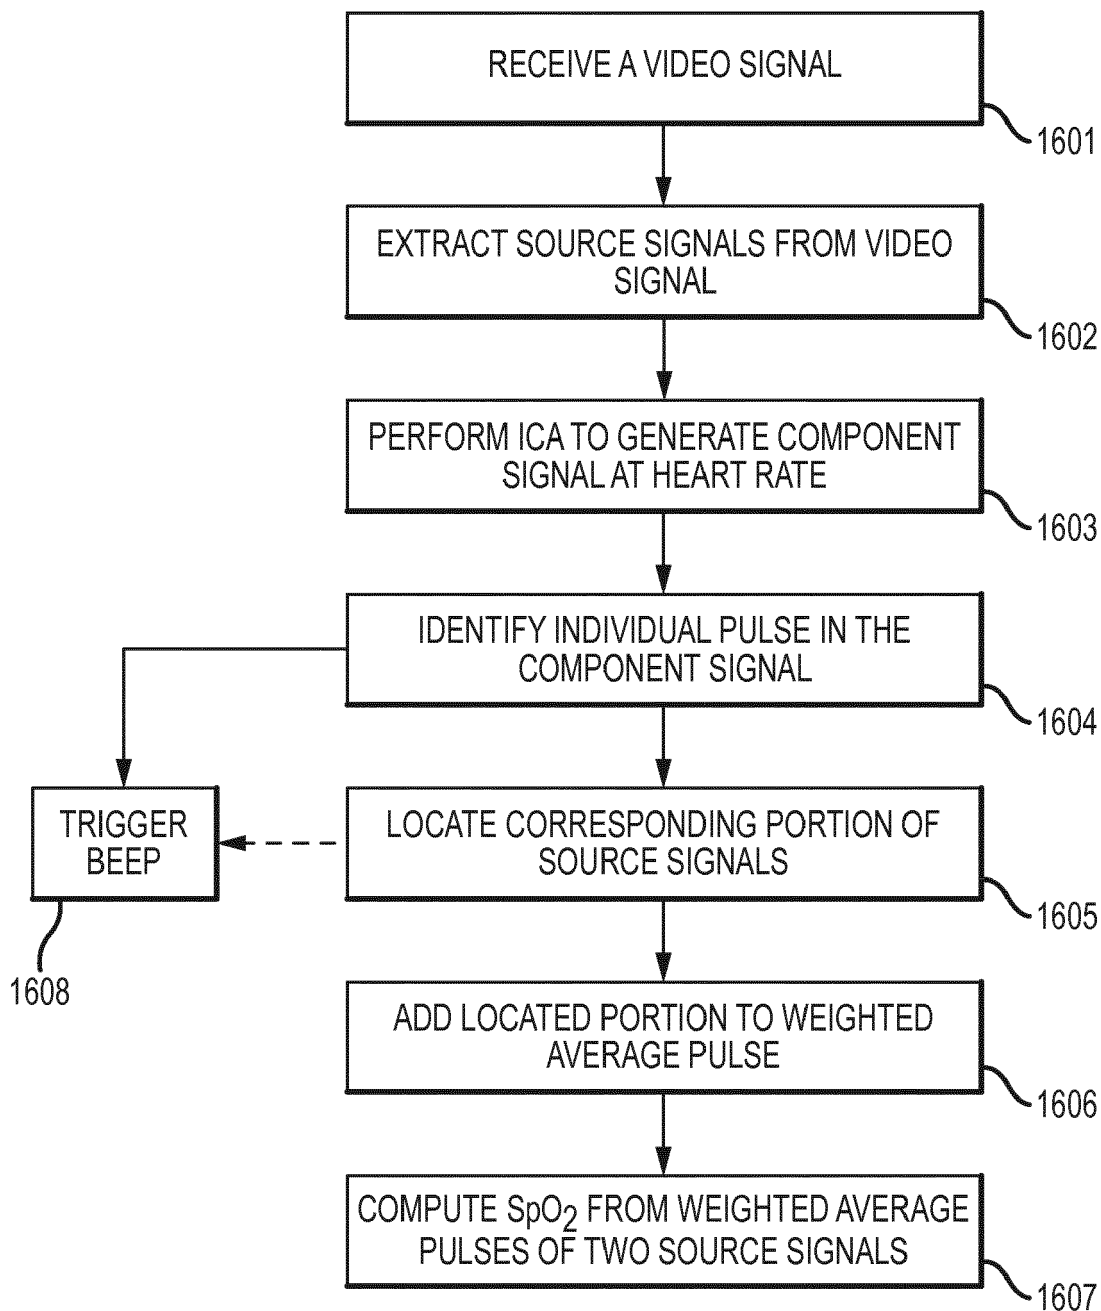


FIG.16

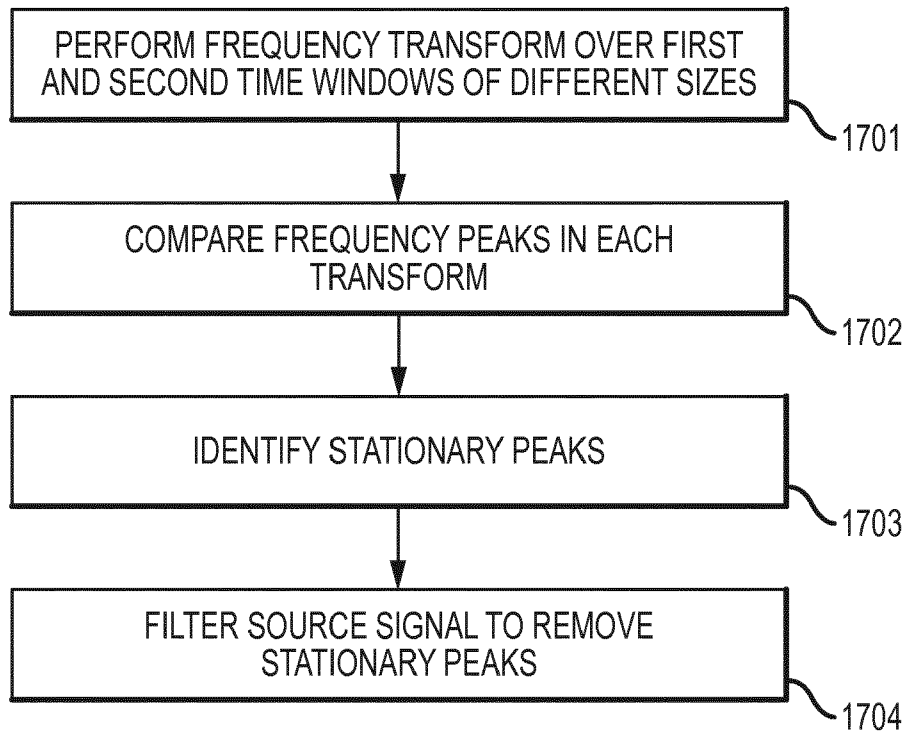


FIG.17A

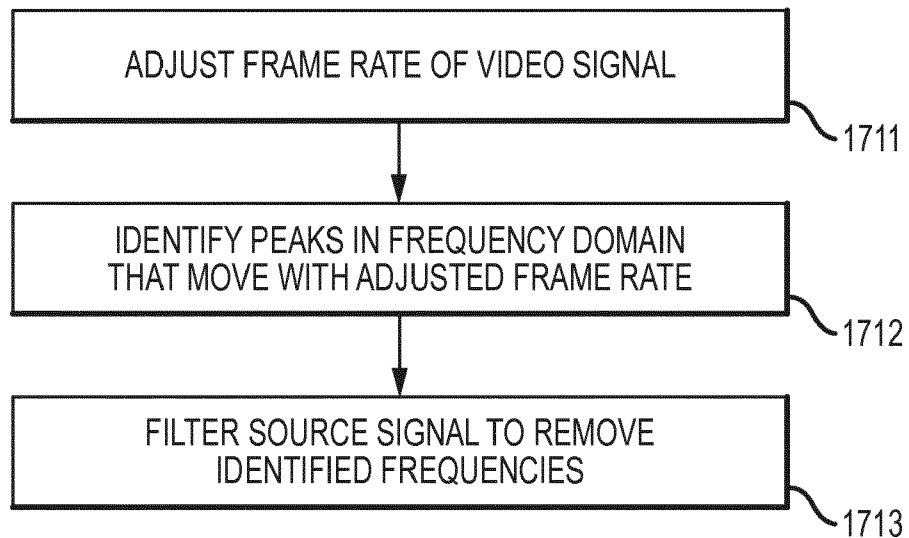



FIG.17B



**FIG. 19**




FIG. 20




FIG. 21



**PARTIAL EUROPEAN SEARCH REPORT**

Application Number

under Rule 62a and/or 63 of the European Patent Convention.  
This report shall be considered, for the purposes of subsequent proceedings, as the European search report

EP 17 15 6334

5

10

15

20

25

30

35

40

45

50

55

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
X	US 2014/023235 A1 (CENNINI GIOVANNI [NL] ET AL) 23 January 2014 (2014-01-23)	1,4	INV. A61B5/00 A61B5/103  ADD. A61B5/024
Y	* paragraphs [0015], [0020] * * paragraph [0053] - paragraph [0058] * * paragraph [0085] - paragraph [0102] * * figures 1,4A-B *	5	
X	US 2014/275832 A1 (MUEHLSTEFF JENS [DE] ET AL) 18 September 2014 (2014-09-18)	1-3,6-9	
Y	* paragraph [0030] * * paragraph [0036] - paragraph [0039] * * paragraph [0079] * * paragraph [0089] - paragraph [0093] * * paragraph [0101] - paragraph [0103] * * paragraph [0108] - paragraph [0109] * * figures 1-7 *	5	TECHNICAL FIELDS SEARCHED (IPC)  A61B
A	WO 2015/078735 A1 (KONINKL PHILIPS NV [NL]) 4 June 2015 (2015-06-04)	1-9	
	* page 9, line 7 - page 11, line 22 * * page 17, line 33 - page 18, line 23 * * figures 1,2,6 *		
----- -/--			
<b>INCOMPLETE SEARCH</b>			
The Search Division considers that the present application, or one or more of its claims, does/do not comply with the EPC so that only a partial search (R.62a, 63) has been carried out.			
Claims searched completely :			
Claims searched incompletely :			
Claims not searched :			
Reason for the limitation of the search: see sheet C			
Place of search		Date of completion of the search	Examiner
The Hague		6 July 2017	Görlach, Tobias
CATEGORY OF CITED DOCUMENTS		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons	
X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document		& : member of the same patent family, corresponding document	

EPO FORM 1503 03.82 (P04E07)



PARTIAL EUROPEAN SEARCH REPORT

Application Number  
EP 17 15 6334

5

10

15

20

25

30

35

40

45

50

55

DOCUMENTS CONSIDERED TO BE RELEVANT			CLASSIFICATION OF THE APPLICATION (IPC)
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	
A	EP 2 772 828 A1 (SONY COMP ENTERTAINMENT INC [JP]) 3 September 2014 (2014-09-03) * paragraph [0016] - paragraph [0026] * * paragraph [0039] - paragraph [0043] * * figures 2,4 *  -----	1-9	
			TECHNICAL FIELDS SEARCHED (IPC)

EPO FORM 1503 08.82 (P04C10) 1

INCOMPLETE SEARCH  
SHEET CApplication Number  
EP 17 15 6334

5

Claim(s) completely searchable:  
1-9

10

Claim(s) not searched:  
10-15

Reason for the limitation of the search:

15

1            Claims 1, 10 and 13 have been drafted as separate independent claims within the same category (method). According to Article 84 in combination with Rule 43(2) EPC, this is only allowable if these claims relate to:

20

- a. a plurality of interrelated products;
- b. alternative uses for the same product; or
- c. alternative solutions to the same problem.

In the present case, however, none of these exceptions applies for the following reasons:

25

1.1            While methods can in principle be interrelated, this is not the case for the methods of claims 1, 10 and 13, because all these methods relate to measuring a vital sign, so that they are in no way complementary or interrelated to each other. Secondly, regarding exception (b), claims 1, 10 and 13 are not drafted as use claims, and it is not apparent from any of these claims that the same product (or apparatus) is being used. Thirdly, regarding exception (c), it is noted that the solutions provided in claims 1, 10 and 13 are neither clearly alternative (which, according to the Guidelines, F-IV, 3.2, last paragraph, is to be interpreted as "different or mutually exclusive solutions"), but are at least partly overlapping (e.g. claim 1 relates to a color signal, which at least partly overlaps with the intensity signal of claim 10; claim 1 relates to a vital sign in general, whereas claim 10 specifically relates to heart rate), and, claim 1 being partly broader ("vital sign") and partly more restricted ("color signal") than claim 10, the underlying technical problems are not the same. This is even more the case for claim 13, which includes receiving a user input and defining one set of pixels (rather than a plurality of regions, as in claims 1 and 10).

30

35

40

1.2            A clarification request under Rule 62a EPC was sent on 28 March 2017. In response to this request, the applicant indicated that claim 1 should be searched without providing further arguments. Consequently, the search of the present application has been limited to claims 1-9.

45

50

55

ANNEX TO THE EUROPEAN SEARCH REPORT  
ON EUROPEAN PATENT APPLICATION NO.

EP 17 15 6334

5 This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.  
The members are as contained in the European Patent Office EDP file on  
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

06-07-2017

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2014023235 A1	23-01-2014	CN 102341811 A	01-02-2012
		EP 2404261 A1	11-01-2012
		JP 5676492 B2	25-02-2015
		JP 2012519527 A	30-08-2012
		US 2011311143 A1	22-12-2011
		US 2014023235 A1	23-01-2014
		WO 2010100593 A1	10-09-2010
US 2014275832 A1	18-09-2014	CN 105188521 A	23-12-2015
		EP 2967377 A1	20-01-2016
		JP 2016513517 A	16-05-2016
		RU 2015143671 A	27-04-2017
		US 2014275832 A1	18-09-2014
		WO 2014140978 A1	18-09-2014
WO 2015078735 A1	04-06-2015	CA 2931377 A1	04-06-2015
		CN 105792742 A	20-07-2016
		EP 3073905 A1	05-10-2016
		JP 2016539697 A	22-12-2016
		US 2017164904 A1	15-06-2017
		WO 2015078735 A1	04-06-2015
EP 2772828 A1	03-09-2014	CN 104054038 A	17-09-2014
		EP 2772828 A1	03-09-2014
		JP 5822651 B2	24-11-2015
		JP 2013090847 A	16-05-2013
		US 2014254902 A1	11-09-2014
		WO 2013061492 A1	02-05-2013

**REFERENCES CITED IN THE DESCRIPTION**

*This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.*

**Patent documents cited in the description**

- US 62297682 A [0001]
- US 62335862 A [0001]
- US 62399741 A [0001]

**Non-patent literature cited in the description**

- **HYVÄRINEN, A. ; OJA, E.** Independent component analysis: algorithms and applications. *Neural networks*, 2000, vol. 13 (4), 411-430 [0113]

专利名称(译)	基于视频的生命体征监测方法		
公开(公告)号	<a href="#">EP3207862A1</a>	公开(公告)日	2017-08-23
申请号	EP2017156334	申请日	2017-02-15
[标]申请(专利权)人(译)	柯惠有限合伙公司		
申请(专利权)人(译)	COVIDIEN LP		
当前申请(专利权)人(译)	COVIDIEN LP		
[标]发明人	ADDISON PAUL STANLEY FOO DAVID JACQUEL DOMINIQUE		
发明人	ADDISON, PAUL STANLEY FOO, DAVID JACQUEL, DOMINIQUE		
IPC分类号	A61B5/00 A61B5/103 A61B5/024		
CPC分类号	A61B5/0077 A61B5/02416 A61B5/1032 A61B5/746 A61B5/7485 A61B5/0205 A61B5/024 A61B5/0816 A61B5/1176 A61B5/14542 A61B5/14551 A61B5/441 A61B5/6843 A61B5/7203 A61B5/7221 A61B5 /7278 A61B5/7425 A61B5/743 A61B5/748 A61B2562/0233 A61B2576/00		
优先权	62/297682 2016-02-19 US 62/335862 2016-05-13 US 62/399741 2016-09-26 US		
外部链接	<a href="#">Espacenet</a>		

摘要(译)

本发明涉及医学监测领域，尤其涉及脉搏率，呼吸率，运动和氧饱和度的非接触式，基于视频的监测。描述了用于捕获患者图像，从图像产生强度信号，过滤那些信号以聚焦在生理成分上，以及从过滤后的信号中测量生命体征的系统和方法。




FIG.2A