



US 20060284979A1

(19) **United States**

(12) **Patent Application Publication**

Clarkson

(10) **Pub. No.: US 2006/0284979 A1**

(43) **Pub. Date: Dec. 21, 2006**

(54) **ACTIVITY RECOGNITION APPARATUS,
METHOD AND PROGRAM**

Publication Classification

(51) **Int. Cl.**
H04N 7/18 (2006.01)

(52) **U.S. Cl.** **348/143**

(75) **Inventor: Brian Clarkson, Tokyo (JP)**

(57) **ABSTRACT**

Correspondence Address:
**C. IRVIN MCCLELLAND
OBLON, SPIVAK, MCCLELLAND, MAIER &
NEUSTADT, P.C.
1940 DUKE STREET
ALEXANDRIA, VA 22314 (US)**

There is provided an activity recognition apparatus for detecting an activity of a subject. The apparatus includes: a sensor unit including a plurality of linear motion sensors configured to detect linear motions and a plurality of rotational motion sensors, the linear motions being orthogonal to each other, the rotational motions being orthogonal to each other; and a computational unit configured to receive and process signals from the sensors included in the sensor unit so as to detect an activity of the subject. The sensor unit is directly or indirectly supported by the subject with an arbitrary orientation with respect to the subject. The computational unit performs a calculation that uses the signals from both linear motion sensors and rotational motion sensors to determine the activity of the subject independent of the orientation of the sensor unit.

(73) **Assignee: Sony Corporation, Tokyo (JP)**

(21) **Appl. No.: 11/435,746**

(22) **Filed: May 18, 2006**

(30) **Foreign Application Priority Data**

Jun. 9, 2005 (JP) 2005-169507

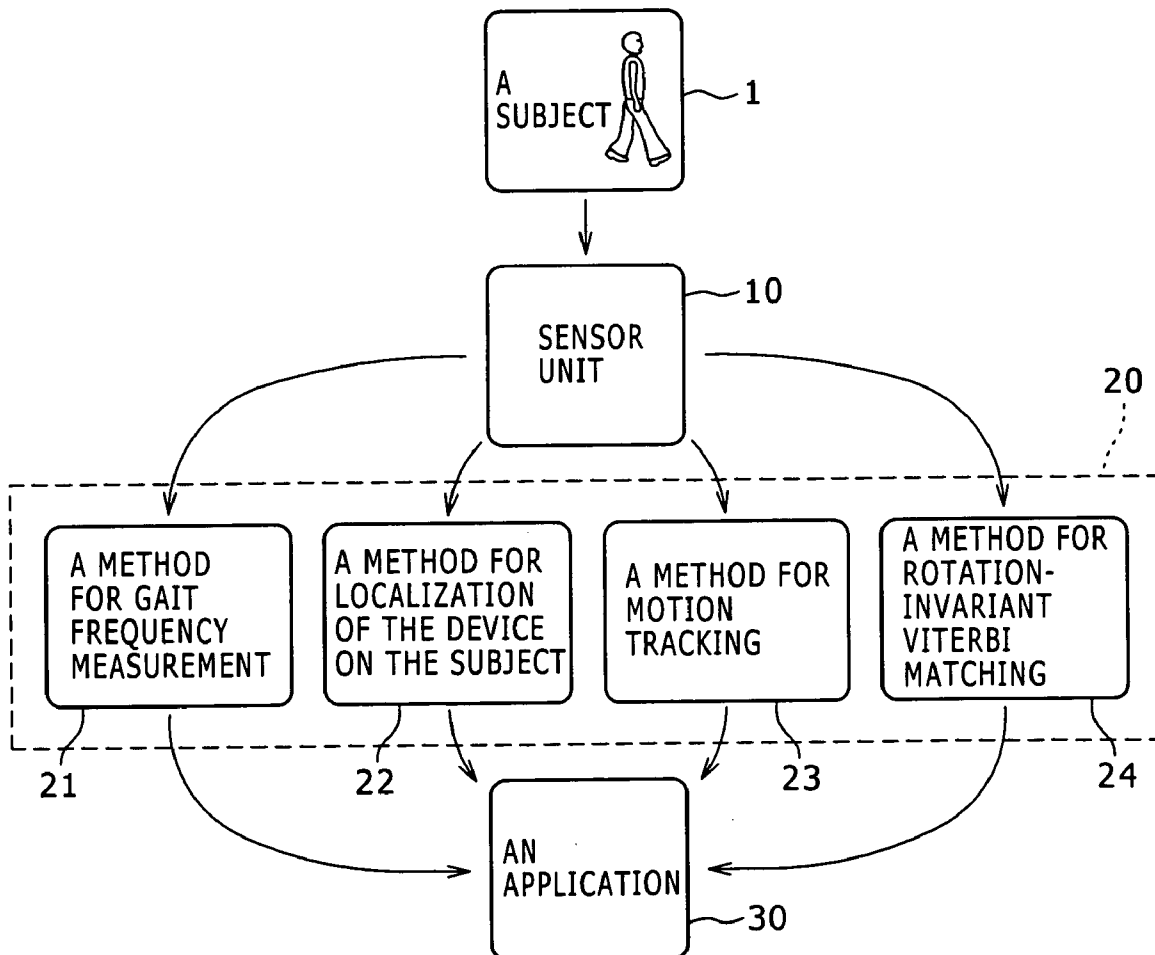


FIG. 1




FIG. 2

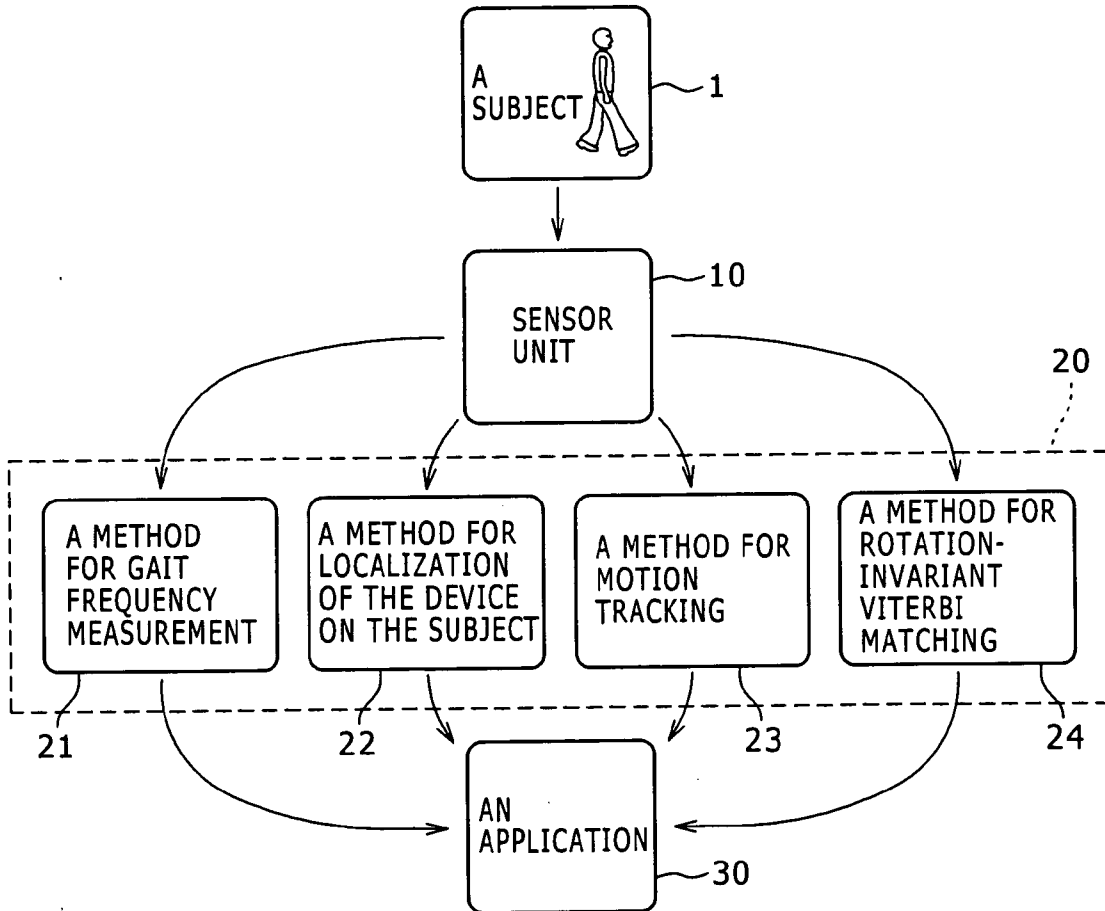


FIG. 3 A




FIG. 3 B




FIG. 4

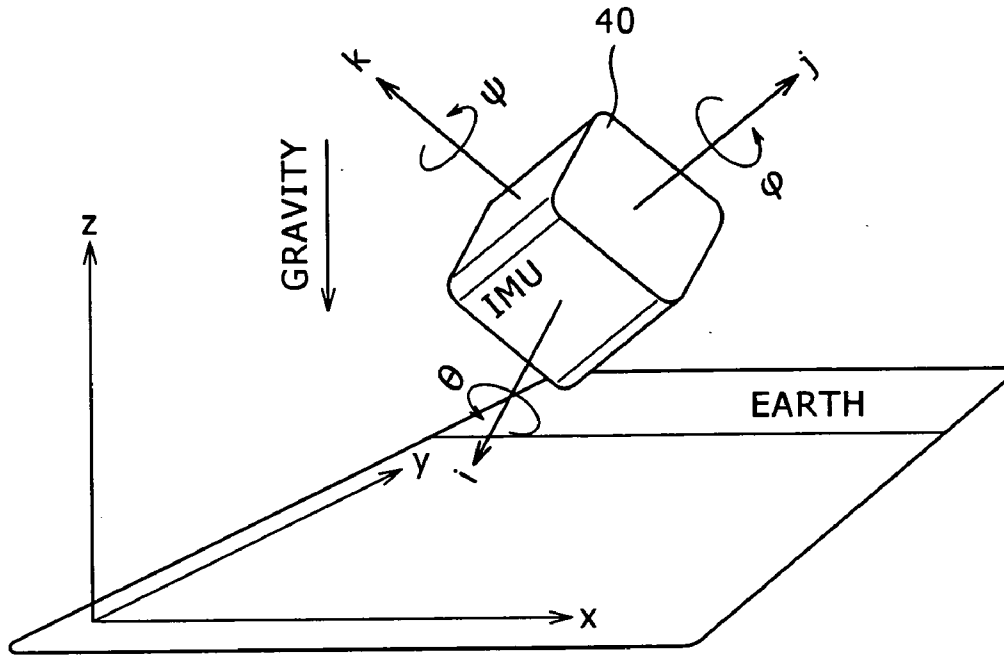


FIG. 5

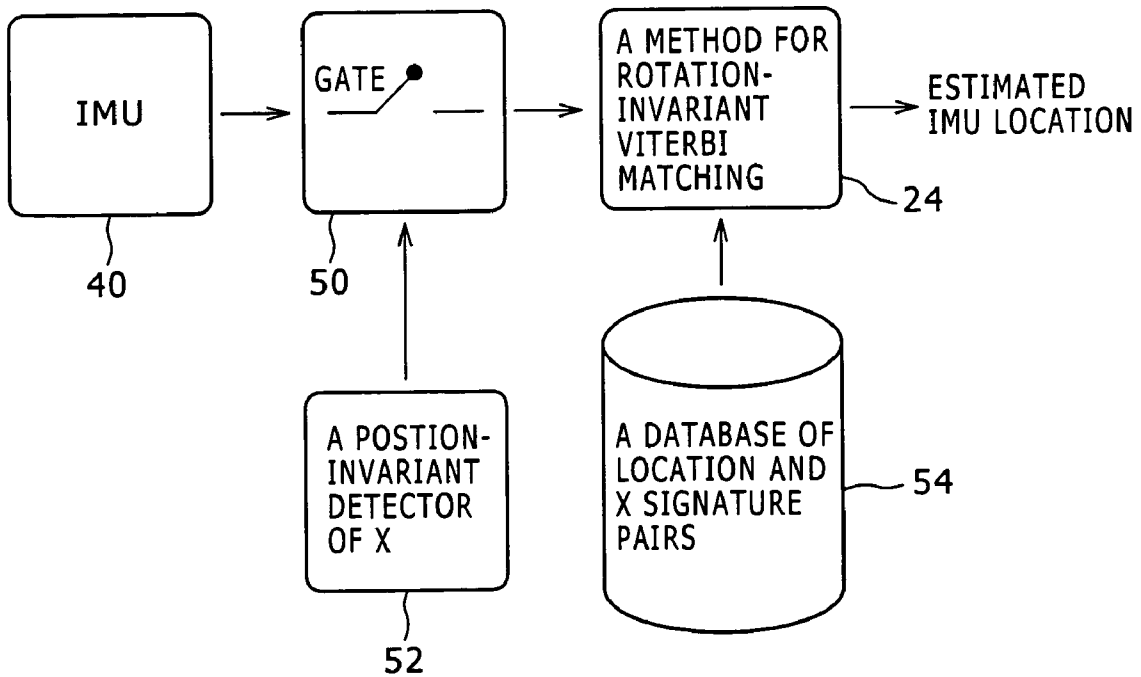


FIG. 6




FIG. 7




FIG. 8 A




FIG. 8 B




FIG. 8 C





FIG. 9



ACTIVITY RECOGNITION APPARATUS, METHOD AND PROGRAM

BACKGROUND OF THE INVENTION

[0001] 1. Field of the Invention

[0002] The present invention relates to an apparatus, method, program and storage medium for activity recognition. More specifically, the present invention relates to an activity recognition apparatus, an activity recognition method, a program to cause a computer to perform such a method, a storage medium or memory storing such a program, and a signal coded to cause a computer to perform such a method.

[0003] 2. Description of Related Art

[0004] Human activity classification has been attempted with cameras, microphones (Non-patent Document 5), inertial sensors (Non-patent Document 4), and goniometers. For each sensor or combination of sensors, researchers have also developed custom algorithms that to varying degrees can classify activity. Since one part of human activity classification is obtaining accurate measurements of human motion, the technology used overlaps with methods used for character animation, kinesiology and biomechanics. However, since the inventor of the present application are addressing the problem of how to determine and classify human activities in a mobile and unconstrained context, all except a few sensor modalities are appropriate.

[0005] Computer Vision

[0006] One set of such constrained methods involves monitoring the subject with static cameras or wearable cameras (Non-patent Document 10). Often optical markers are attached to the subject's body to simplify visual tracking. The subject's movement can be inferred from the movement of the markers through an inverse-kinematic model.

[0007] While this method has the potential of providing very rich and complete measurements, it can only be used in the well-lit space in view of the cameras, which precludes its use in mobile applications.

[0008] Computer Audition

[0009] There also examples where researchers have tried using microphones (both on the body and in the environment) to automatically determine a user's activity (see Non-patent Document 2 and Non-patent Document 9). However, major advantages of motion sensors are lower data rate (which translates to much lower computational requirements) and the ability to measure only the motion of the user. In audio, sounds caused by the user's activities are mixed with an ever-changing background of environmental sounds. Separation of these foreground and background is an unsolved problem which does not arise when using motion sensors.

[0010] Absolute Joint Angle Sensing

[0011] This set of methods involves attaching sensors directly to the subject's body so as to measure the angles of the subject's joint. Possible sensor's are potentiometers and rotational encoders (which both require an exoskeleton to transmit the changes in a subject's joint angles to the sensor), bend sensors (strips of material that are sensitive to bending and flexing), and inertial sensors (accelerometers

and gyros, which essentially measure changes in velocity via Newton's First Law of Motion). In clinical situations, goniometers (equivalent to a tilt sensor) are often used to measure angles with respect to gravity.

[0012] Inertial Sensing

[0013] Of all these only the inertial sensors and tilt sensors can be used to measurement of a subject's movements without requiring extensive modifications to clothing (for exoskeletons or strips of flex-sensitive material) or requiring special attachments. Furthermore, inertial sensors can be completely enclosed inside another device that user can simply hold or carry in a pocket or bag, thus making these sensors attractive for use in mobile devices or as self-contained devices.

[0014] Other relevant works are disclosed, for example, in Non-patent Document 7, Non-patent Document 8, Patent Document 1, Patent Document 2.

REFERENCES

[0015] Non-Patent Document 1: Dimitri P. Bertsekas. *Nonlinear Programming*. Athena Scientific, 2nd edition, 1995.

[0016] Non-Patent Document 2: Brian Clarkson, Alex Pentland, and Kenji Mase. Recognizing user context via wearable sensors. In *Fourth International Symposium on Wearable Computers (ISWC)*, 2000.

[0017] Non-Patent Document 3: Masataka Ishikawa. U.S. Pat. No. 5,297,028: Method and apparatus for correcting drift errors in an angular rate sensor. Filed by Zexel Corporation Daihatsu-Nissan, Tokyo, Japan, August 1991.

[0018] Non-Patent Document 4: Holger Junker, Jamie Ward, Paul Lukowicz, and Gerhard Troster. Continuous recognition of arm activities with body-worn inertial sensors. In *Proceedings of the 8th International Symposium on Wearable Computers*, 2004.

[0019] Non-Patent Document 5: Paul Lukowicz, Jamie Ward, Holger Junker, Mathias Stager, Gerhard Troster, Amin Atrash, and Thad Starner. Recognizing workshop activity using body-worn microphones and accelerometers. In *Second IEEE International Conference on Pervasive Computing and Communications*, 2004.

[0020] Non-Patent Document 6: Thomas P. Minka. Old and new matrix algebra useful for statistics. Technical report, Massachusetts Institute of Technology, 2000.

[0021] Non-Patent Document 7: Bruno Scherzinger. U.S. Pat. No. 6,853,909: Walking stick navigator for position determination. Filed by Applanix Corporation, Inc., Ontario, Canada, November 2002.

[0022] Non-Patent Document 8: Douglas Smith, Aaron Joseph, David Boone, Robert Borchers, and Ernest Burgess. U.S. Pat. No. 5,485,402: Gait activity monitor. Filed by Prosthetics Research Study, Seattle, Wash., March 1994.

[0023] Non-Patent Document 9: Mathias Staeger, Paul Lukowicz, and Gerhard Troester. Implementation and evaluation of a low-power sound-based user activity recognition system. In *Eighth International Symposium on Wearable Computers (ISWC)*, 2004.

[0024] Non-Patent Document 10: Thad Starner, Josh Weaver, and Alex Pentland. Real-time american sign language recognition using desk and wearable computer-based video. IEEE Transactions on Pattern Analysis and Machine Intelligence, 20(12):1371-175, December 1998.

[0025] Non-Patent Document 11: Tero Tolonen and Matti Kajjalainen. A computationally efficient multipitch analysis model. IEEE Transactions on Speech and Audio Processing, 8(6), November 2000.

[0026] Patent Document 1: Robert Levi and Thomas Judd. U.S. Pat. No. 5,583,776: Dead reckoning navigational system using accelerometer to measure foot impacts. Filed by Point Research Corporation, Santa Ana, Calif., March 1995.

[0027] Patent Document 2: Christopher Townsend and Steven Arms. U.S. Pat. No. 6,834,436: Posture and body movement system. Filed by MicroStrain, Inc., Williston, Vt., February 2002.

SUMMARY OF THE INVENTION

[0028] However, obtaining accurate motion measurements with inertial sensors has historically proven to be quite difficult due to the sensitivity these sensors have to shock, temperature, electrical noise and so on. For example, gyros (angular rate sensors) have a dependency on temperature which must be accounted for Non-patent Document 3.

[0029] Calibration is also required since these devices will typically have an arbitrary offset and scale factor affecting their outputs.

[0030] Further, inertial sensors have no notion of the subject's frame of reference and measurements are always relative to the physical sensor. One exception to this, is that a properly oriented accelerometer can measure the earth's gravitational field and thus providing a way to determine the down direction. The accelerometer must not be undergoing any other unknown accelerations for this to work.

[0031] Accordingly, in view of the above-mentioned issues, it is desirable to perform activity recognition using inertial sensors without posing any preconditions on orientation of the inertial sensors relative to the subject's reference frame, thereby allowing more flexibility in implementing a system for the activity recognition.

[0032] According to an embodiment of the present invention, there is provided an activity recognition apparatus for determining an activity of a subject. The activity recognition apparatus includes: a sensor unit including a plurality of linear motion sensors configured to detect linear motions and a plurality of rotational motion sensors, the linear motions being orthogonal to each other, the rotational motions being orthogonal to each other; and a computational unit configured to receive and process signals from the sensors included in the sensor unit so as to determine an activity of the subject. The sensor unit is directly or indirectly supported by the subject with an arbitrary orientation with respect to the subject; and the computational unit performs processing that uses the signals from both linear motion sensors and rotational motion sensors to determine the activity independent of the orientation of the sensor unit.

[0033] According to another embodiment of the present invention, there is provided an activity recognition method of determining an activity of a subject. The method includes:

sampling accelerations in three orthogonal axes and angular velocities around the three orthogonal axes; and determining an activity of the subject using all the sampled accelerations and angular velocities. In the method, the accelerations and angular velocities are detected by sensors integrated into a single unit, which is directly or indirectly supported by the subject with an arbitrary orientation with respect to the subject; and the determination of activity is performed by using both the accelerations and angular velocities to determine the activity independent of the orientation of the sensor unit.

[0034] According to another embodiment of the present invention, there is provided a program that causes a computer to perform activity recognition by determining an activity of a subject, or a recoding medium that store such a program, or a signal coded to cause a computer to perform such a program. The program includes: receiving data of accelerations in three orthogonal axes and angular velocities around the three orthogonal axes; and determining an activity of the subject using all the received data of accelerations and angular velocities. The accelerations and angular velocities are detected by sensors integrated into a single unit, which is directly or indirectly supported by the subject with an arbitrary orientation with respect to the subject; and the determination of activity is performed by using both the accelerations and angular velocities to determine the activity independent of the orientation of the sensor unit.

[0035] In the embodiments of the present invention, the sensor unit is directly or indirectly supported by the subject with arbitrary orientation. In other words, the signals from the sensor unit do not have any information about the sensor's orientation with respect to the coordinates of the subject. However, different movements of the entire sensor unit correspond with different activities of the subject.

[0036] Accordingly, the activity of the subject can be determined independent of the sensor unit's orientation by detecting the movement of the entire sensor unit and by associating the detected movement with a corresponding activity of the subject or a corresponding motion signature that associated with the activity.

[0037] According to the present invention, it is possible to perform activity recognition using inertial sensors without posing any preconditions on orientation of inertial sensors relative to the subject's reference frame, thereby allowing more flexibility in implementing a system for the activity recognition.

BRIEF DESCRIPTION OF THE DRAWINGS

[0038] The above and other objects, features and advantages of the present invention will become more apparent from the following description of the presently preferred exemplary embodiments of the invention taken in conjunction with the accompanying drawings, in which:

[0039] FIG. 1 is a block diagram of an activity recognition apparatus according to an embodiment of the present invention;

[0040] FIG. 2 is a conceptual overview of an activity recognition apparatus according to an embodiment of the present invention;

[0041] FIGS. 3(a) and 3(b) are schematic diagrams of activity recognition apparatuses for different usage scenarios;

[0042] FIG. 4 is an conceptual overview of a physical module (IMU) according to an embodiment of the present invention;

[0043] FIG. 5 is an overview of processing pipeline for a method of determining a location of an IMU on a user's body when using a reference motion X;

[0044] FIG. 6 is an overview of processing pipeline for a method of tracking an IMU's motion;

[0045] FIG. 7 is a schematic diagram showing a geometric construction used to directly calculate sine and cosine of a half-angles required by quaternion representation in a Plumb model;

[0046] FIG. 8(a) and FIG. 8(b) show example sequences of motion signatures, and FIG. 8(c) shows matching of two sequences according to an embodiment of the present invention; and

[0047] FIG. 9 is a schematic diagram of a computational unit for a method of comparing motion signatures in accordance with an embodiment of the present invention.

DETAILED DESCRIPTION OF THE EMBODIMENTS

[0048] Embodiments of the present invention are described with reference to accompanying figures.

[0049] FIG. 1 shows a configuration example of an activity recognition apparatus for determining an activity of a subject in accordance with an embodiment of the present invention. The apparatus includes a sensor unit 10 and a computational unit 20. The sensor unit 10 includes a plurality of linear motion sensors configured to detect linear motions and a plurality of rotational motion sensors, the linear motions being orthogonal to each other, the rotational motions being orthogonal to each other. The computational unit 20 is configured to receive and process signals from the sensors included in the sensor unit so as to output an activity of the subject. The sensor unit 10 is directly or indirectly supported by the subject with an arbitrary orientation with respect to the subject. The computational unit 20 performs a calculation that uses the signals from both linear motion sensors and rotational motion sensors to determine the activity independent of the orientation of the sensor unit.

[0050] The plurality of linear motion sensors may be three accelerometers for measuring accelerations in three orthogonal axes, and the plurality of rotational motion sensors may be three gyro sensors for measuring angular velocities around the three orthogonal axes. The computational unit 20 may use six outputs from the accelerometers and the gyro sensors to determine a current activity of the subject.

[0051] The computational unit may includes; a buffer configured to store an observation sequence formed of a plurality set of the sensor outputs for a predetermined period of time; a storage configured to store a plurality of reference sequences, the reference sequences corresponding to motion signatures of different activities; and a matching processor configured to match the observation sequence against the reference sequences to find a best matched reference sequence. Further, the matching processor may find an optimal rotation and temporal correspondence between the observation sequence and one of the reference sequences so as to obtain the best matched reference sequence.

[0052] Furthermore, the reference sequence may include a plurality of vectors representing states in Hidden Markov Model; and the matching between the observation sequence and the reference sequence may be performed using Viterbi algorithm so as to obtain an optimal state path.

[0053] Alternatively, the computational unit may further detect a rate of the subject's gait by calculating an autocorrelation of the signals from the linear motion sensors and an autocorrelation of the signals from the rotation motion sensors, and summing these autocorrelations to obtain a summary autocorrelation function; and the summary autocorrelation function may be used to determine the rate of the subject's gait.

[0054] Alternatively, the computational unit may further detect a rotation of the sensor unit by calculating a first rotation by integrating the signals from the rotational motion sensors, calculating a second rotation by using a gravitation model and the signals from the linear motion sensors, and performing weighted interpolation of the first and second rotation. Further, the weight for the first rotation is larger if variations in the signals from the rotational motion sensors are smaller; and the weight for the second rotation is larger if variations in the signals from the linear motion sensors are smaller.

[0055] Alternatively, the activity recognition apparatus may further include a storage configure to store reference motion signatures and corresponding locations for a reference activity. Furthermore, the computational unit may further detect a location of the sensor unit on the subject by determining if a current activity represented with the signals from the sensors is the reference activity, and, if the current activity is the reference activity, matching a motion signature corresponding to the current activity against the reference motion signature so as to determine a location of the sensor unit, which corresponds to a best matched reference motion signature.

[0056] According to another embodiment of the present invention, there is provided an activity recognition method of determining an activity of a subject. The method includes: sampling accelerations in three orthogonal axes and angular velocities around the three orthogonal axes; and determining an activity of the subject using all the sampled accelerations and angular velocities. Furthermore, the accelerations and angular velocities are detected by sensors integrated into a single unit, which is directly or indirectly supported by the subject with an arbitrary orientation with respect to the subject; and the determination of activity is performed by using both the accelerations and angular velocities to determine the activity independent of the orientation of the sensor unit.

[0057] According to another embodiment of the present invention, there is provided a program that causes a computer to perform activity recognition by determining an activity of a subject. The program includes: receiving data of accelerations in three orthogonal axes and angular velocities around the three orthogonal axes; and determining an activity of the subject using all the received data of accelerations and angular velocities. Furthermore, the accelerations and angular velocities are detected by sensors integrated into a single unit, which is directly or indirectly supported by the subject with an arbitrary orientation with respect to the subject; and the determination of activity is performed by

using both the accelerations and angular velocities to determine the activity independent of the orientation of the sensor unit.

[0058] Below, still another embodiment according to the present invention is described. In the following embodiment, there is provided a system that is capable of activity recognition and includes a sensor unit and a computational unit. In the following embodiment, it is assumed that the sensor unit is “supported” by a human subject and the computational unit determines a type of current activity of the human subject based on the outputs from the sensor unit.

[0059] (1) System Overview

[0060] In the present embodiment, as shown in FIG. 2, a six degree-of-freedom (DOF) inertial sensor is employed as the sensor unit 10 to classify various common human activities in real-time. In the following description, the term “the Sensor” indicates the 6-DOF inertial sensor, and “the User” denotes the human subject 1 who is either carrying or wearing the Sensor. Furthermore, the term “the Device” denotes the present system combining an inertial measurement unit (IMU) with the Sensor and the above-mentioned computational unit 20 for performing steps of methods for classification of the User’s activity. Results of classification may be outputted or transfer to an application.

[0061] While using inertial sensors to measure human activity has been attempted before, the present system defining characteristic is its ability to operate fully automatically almost anywhere on the human body with any sensor orientation. In the present embodiment, the system incorporates four distinct methods for different aspects of the User’s activity and device state, each of which can be performed without posing any preconditions on the sensor orientation. The methods includes a method 21 for gait frequency measurement, a method 22 for localization of the Device on the User, a method 23 for motion tracking, and a method 24 for rotation-invariant Viterbi matching (see FIG. 2).

[0062] The Sensor may be supported or held by the User directly or indirectly. For example, the following (nonexhaustive) manner of coupling may be possible:

[0063] The User is holding the Sensor in his hands.

[0064] The User is carrying the Sensor in a purse or backpack which is being worn.

[0065] The Sensor is rigidly or semi-rigidly attached to the User’s body or clothing.

[0066] The Sensor is hanging from the User’s neck or waist.

[0067] The Sensor has been integrated into the construction of the User’s shoes, jewelry, hat, gloves, portable electronic device or any other object whose normal use involves being worn or carried by the User.

[0068] The Sensor is attached or place inside a vehicle that is simultaneously transporting the User.

[0069] The present embodiment may also be applicable even in cases where the Sensor is coupled to the User with any arbitrary forms as long as the Sensor’s motions have a relationship with the User’s motions.

[0070] (2) Physical Device

[0071] For the purposes of a simple and concrete description of the present embodiment, it is assumed that the 6-DOF inertial sensor is composed of components that measure linear acceleration in the x-, y-, and z-directions and angular velocity around the x-, y-, and z-axis. However, the methods described hereafter can be trivially adapted to other schema for measuring 6-DOF.

[0072] In the present embodiment, as shown in FIG. 3(a), the system may include a sensor mote 10a and a computational unit 20 embedded in a host device 30a. The sensor mote 10a includes a 6-DOF inertial sensor 10a-1 and a wireless transmitter 10a-2 for transmitting sensor data to the computational unit 20a. Alternatively, a memory may be provided instead of the wireless transmitter 10a-2 for storing sensor data for a certain period of time. The stored data may be transferred to the computational unit 20a when the sensor mote 10a is connected to the host device 30a.

[0073] Alternatively, the entire system including a 6g-DOF inertial sensor 10b and a computational unit 20b may be embedded in a single host device 30b such as a cellphone or portable game apparatus as shown in FIG. 3(b).

[0074] (3) Sensor Signal Conditioning and Calibration

[0075] Below, there is described the computational processing of the raw sensor data from the Sensor so as to yield classifications of the User’s activities in real-time. First, some symbols representing the raw output of the Sensor are defined:

$$u(t) = \begin{bmatrix} s_x(t) \\ s_y(t) \\ s_z(t) \end{bmatrix} \in R^3 \quad (1)$$

$$v(t) = \begin{bmatrix} s_\theta(t) \\ s_\phi(t) \\ s_\psi(t) \end{bmatrix} \in R^3 \quad (2)$$

[0076] Each sample from the Device’s sensors contains the six values specified in Equations (1) and (2). The $u(t)$ sensor values are linearly proportional to the acceleration of the Device along the i-, j-, and k-axes. These could be considered to be measured, for example, by MEMS accelerometer or similar devices. The $v(t)$ sensor values are linearly proportional to the angular velocity of the Device around the i-, j-, and k-axes. These 10 could be considered to be obtained, for example, from MEMS gyros or similar devices. FIG. 4 shows all the measurements are made with respect to basis directions in the local coordinate system (i-j-k) of the Device itself.

[0077] If it is supposed that $p(t) \in R^3$ is the position of the Device in any world coordinate system (e.g. x-y-z in FIG. 4), $g(t) \in R^3$ is the acceleration of the Earth’s gravitational field in the local coordinate system, and $q(t) \in R^3$ is the orientation of the Device in its local coordinate system, then it can be written:

$$u(t) = A[\nabla^2 p(t) + g(t)] + u_0 + \epsilon_u(t) \quad (3)$$

$$v(t) = B\nabla q(t) + v_0 + \epsilon_v(t) \quad (4)$$

[0078] All gradients are taken with respect to the local coordinate system of the Device. A and B are 3-by-3 scaling matrices. The diagonal elements specify the scale-factors (also known as sensitivities) which convert the sensor output (e.g. a voltage) to standard physical units (e.g. m/s² and °/s). At the same time, the off-diagonal elements describe the cross-talk between sensors due to slight misalignments (for example, the angle between the θ and ϕ gyros is not exactly 90°). The arbitrary offsets (also known as bias) exhibited by typical sensors are denoted by u_0 and v_0 . The sensor noise, which includes the quantization noise, is specified for each sensor as a sequence of random variables, $\epsilon_u(t)$ and $\epsilon_v(t)$.

[0079] Equation (3) and (4) describe how the actual motion of the Device translates into sensor measurements. If it is solved for acceleration and angular velocity using the expressions in (3) and (4), then the following expressions are obtained:

$$\nabla^2 p(t) = A^{-3} [u(t) - u_0 - \epsilon_u(t)] - g(t) \quad (5)$$

$$\nabla q(t) = B^{-1} [v(t) - v_0 - \epsilon_v(t)] \quad (6)$$

[0080] Signal calibration involves determining the values of A, u_0 , B, v_0 . Depending on the 10 actual components used inside the IMU, these values can be dependent on manufacturing conditions, ambient temperature, driving voltage, and so on. However, only the method for motion tracking require these parameters to be known. All of the other methods 21-22 and 24 given here can use uncalibrated sensor values.

[0081] Finally, either the continuous signals from the IMU are sampled at a given rate and passed onto the computational unit or the signals leave the IMU already digitized. For the remainder of this description, let r_{imu} be the digital sampling rate of the output of the IMU. Thus the digitized output of the IMU can be written as:

$$u[t] = u(t/r_{imu}) \quad (7)$$

$$v[t] = v(t/r_{imu}) \quad (8)$$

$$t = 0, 1, 2, \dots \quad (9)$$

[0082] (4) Method for Gait Frequency Measurement

[0083] This Section describes our method for determining the User's activity, such as determining if the User is walking or running or not moving, and tracking the rate of his/her gait. This method is effective for any location of the device on the User's body. This method for gait detection and tracking involves tracking the period of repetitive movements apparent in the output of the acceleration sensors.

[0084] A common method for detecting and measuring the periods of periodic signals is by finding peaks in the autocorrelation of the signal. To see how the autocorrelation of the acceleration sensor output, u , relates to the autocorrelation of the actual acceleration of the IMU, $\nabla^2 p$, (multivariate) autocorrelation of the i -th frame may be written as:

$$\begin{aligned} u \star u &= \sum_{i=0}^{N_{gait}} u[i] \cdot u[i+t] \\ &= \phi_1 + \phi_2 + \phi_3 \end{aligned} \quad (10)$$

[0085] where \star denotes the cross-correlation operator and,

$$\phi_1 = (A \nabla^2 p \star A \nabla^2 p) + (Ag \star Ag)$$

$$\phi_2 = 2(A \nabla^2 p \star Ag)$$

$$\phi_3 = 2(A \nabla^2 p \star \epsilon_u + A \nabla^2 p \star u_0) + 2(u_0 \star \epsilon_u) + (u_0 \star u_0) + (\epsilon_u \star \epsilon_u)$$

[0086] The first term of Equation (10), ϕ_1 , is the most important term containing the periodicities of the acceleration of the IMU and gravity (which depends only on the orientation of the IMU). The second term, ϕ_2 , is a cross-correlation between acceleration and gravity. For any kind of natural repetitive motion involving the human body, it is pretty much guaranteed that $g(t)$ and $\nabla^2 p$ will share the same periodicities and hence their autocorrelations, ϕ_1 , and cross-correlations, ϕ_2 , will add constructively. The final term, ϕ_3 , contains correlations of the noise and the offset with various terms. These correlations will not exhibit any periodicities and will simply add noise and a constant bias to the final autocorrelation of Equation (10). The validity of the autocorrelation of the gyro sensor outputs, v , follows exactly the same argument.

[0087] In practice, a small time interval is chosen over which to conduct the autocorrelation analysis and a rate is chosen at which to perform the analysis. Let T_{gait} be the length of time for one gait analysis frame (e.g. interval of sensor data) and let r_{gait} be the rate of gait analysis. Thus, in Equation (10),

$$N_{gait} = [T_{gait} r_{IMU}] \quad (11)$$

[0088] is the number of vectors in one analysis frame.

[0089] The present method involves calculating the autocorrelation of u and v and summing the result to get a summary autocorrelation function (SACF) of the total motion of the IMU:

$$p[t] = (u \star u)[t] + (v \star v)[t] \quad (12)$$

[0090] For efficiency and improvement of temporal resolution, the Wiener-Khinchin Theorem is used to calculate a windowed autocorrelation. In general, if $x[t]$ is any sequence, $X[\omega] = F\{x\}$ is its Discrete Fourier Transform (DFT), and $W[\omega] = F\{w\}$ is the DFT of a windowing function (such as a Hamming window), then:

$$x \star x = F\{W \star X / W^2\} \quad (13)$$

[0091] Using the Fast Fourier Transform (FFT) to calculate the DFT in Equation (13) results in a very efficient calculation of the autocorrelations in Equation (12).

[0092] The peaks of the ACF correspond to the periods of any cyclic motion in the data. However, since the ACF will also contain peaks at integer multiples of the true period of any cycles, a peak pruning procedure similar to that introduced in Non-patent Document 11 is used to calculate an enhanced summary autocorrelation function (ESACF). This peak pruning technique is as follows:

[0093] PRUNE-HARMONICS(SACF) returns ESACF

[0094] 1. Clip the SACF to positive values.

[0095] 2. Upsample and interpolate by a factor of two.

[0096] 3. Subtract the upsampled version from the clipped original.

[0097] 4. Clip to positive values again.

[0098] 5. Repeat 2-4 for prime factors three, five, seven, etc.

[0099] The goal of this procedure is to sequentially remove all multiples of each peak, thus leaving only the fundamental periods.

[0100] Now if the highest peak (for non-zero lags) and threshold are chosen, it is possible to determine when there is cyclic motion (such as walking) and at what pitch it is occurring at. If τ^* is the lag of the highest peak in the ESACF, then the pitch of the gait is:

$$P_{gait} = \frac{r_{IMU}}{\tau^*} \quad (14)$$

[0101] If there is no peak with a value over a certain threshold (defined as a percentage of the energy in the signal during the analysis interval), then the system can generate a report that the User is not walking or running (gait pitch=0 Hz). Since, the gait pitch changes over time, the above analysis is done repeatedly for each successive analysis interval, yielding gait pitch estimates at the rate of r_{gait} .

[0102] (5) Method for Motion Tracking

[0103] In this Section, the method 23 for tracking the down direction and heading continuously no matter how the orientation of the device changes or moves is provided. This method assumes that the output from the IMU has been calibrated.

[0104] The 3-axis of gyro sensing data contains enough information to track the IMU's orientation. However, all gyros suffer from an accumulation of drift and other types of error. Thus, using just gyro information to track orientation, while quite accurate in the short-term, leads to large tracking errors in the long term.

[0105] The 3-axis of acceleration sensing data contains information about what direction Earth's gravity is pulling. Knowledge of the direction of the Earth's gravitational pull can be used to track the IMU's orientation drift-free with at most two degrees of freedom. The problem with just using the accelerometers in this fashion is that when the IMU is being acted on (moved by some outside force) the estimate of the Earth's gravitational pull is corrupted.

[0106] It is very difficult and unnatural for humans to move or move another object with constant acceleration or velocity. Similarly, with respect to rotation, it is very difficult for humans to rotate or rotate another object with constant angular velocity.

[0107] In order to robustly estimate the down direction, it is preferable to rely on the acceleration measurements more when the accelerometer signal is constant because it is highly likely that a human is not acting on the device. Conversely, it is preferable to rely more on the gyro measurements when the accelerometer signal is non-constant.

[0108] This principle also applies to online re-zeroing of the gyro signal. In other words, it is preferable that Gyro measurements is re-zeroed when the gyro measurement is constant. Luckily, this is also true when relying mostly on the acceleration measurements for tracking.

[0109] FIG. 6 shows the computational pipeline for the motion tracking method. Linear acceleration data 410 from the IMU 40 including the sensor is used to calculate the

confidence 412 that only gravity is acting on the IMU 40 and as input into the Plumb Model 414 (explained later). Angular velocity data 420 from the IMU 40 is used to calculate the confidence 422 that the IMU 40 is not being rotated and as input into the Gyro Integration module 424.

[0110] Both the Plumb Model 414 and the Gyro Integration 424 output separate estimates of the IMU's orientation, i.e. Q_{plumb} and Q_{gyro} (described later). Q_{plumb} and Q_{gyro} are then combined using the confidence scores to weight one versus the other (430) to perform weighted interpolation to calculate the estimated IMU orientation. The Gyro Integration computes an orientation relative to the orientation of the previous timestep and thus requires feedback.

[0111] The Plumb Model 414 gives an absolute estimate of orientation and thus requires no feedback.

[0112] (5.1) The Plumb Model

[0113] The Plumb Model uses the Earth's gravitational pull to measure orientation, albeit only with two degrees of freedom. In short, this module outputs the rotation that maps the current down vector to a unit reference vector, r , in the coordinate system of the IMU 40 (see FIG. 4). This reference vector can be chosen arbitrarily, or during zeroing. The resulting rotation is given in quaternion form as:

$$Q_{plumb}[t] = \begin{bmatrix} a_x[t]\sin(\theta[t]/2) \\ a_y[t]\sin(\theta[t]/2) \\ a_z[t]\sin(\theta[t]/2) \\ \cos(\theta[t]/2) \end{bmatrix} \quad (15)$$

[0114] where $a[t] = (a_x[t], a_y[t], a_z[t])^T$ is the axis of rotation and $\theta[t]$ is the angle of rotation. The down vector is just the normalized linear acceleration (assuming no forces other than gravity are acting on the IMU 40):

$$d[t] = \frac{u[t]}{\sqrt{u[t]^T u[t]}} \quad (16)$$

[0115] The cross product of the reference vector and the down vector gives the axis of rotation:

$$a[t] = d[t] \times r \quad (17)$$

[0116] The $\sin(\theta[t]/2)$ and $\cos(\theta[t]/2)$ are calculated directly (without using inverse trigonometric functions which are hard to implement in embedded applications) from FIG. 7 as:

$$\sin(\theta[t]/2) = \frac{|r + d[t]|}{2} \quad (18)$$

$$\cos(\theta[t]/2) = \frac{|r - d[t]|}{2} \quad (19)$$

[0117] (5.2) Gyro Integration

[0118] Integration of the gyro's angular velocity measurements is straightforward. As input, the previous orientation

of the IMU, $Q[t-1]$ is required. First, each axial rotation is represented as a quaternion rotation in local coordinates (i-j-k in **FIG. 4**):

$$q_i = (\cos(s_{\theta}[t]/2), (1, 0, 0)^T \sin(s_{\theta}[t]/2)) \quad (20)$$

$$q_j = (\cos(s_{\phi}[t]/2), (0, 1, 0)^T \sin(s_{\phi}[t]/2)) \quad (21)$$

$$q_k = (\cos(s_{\psi}[t]/2), (0, 0, 1)^T \sin(s_{\psi}[t]/2)) \quad (22)$$

[0119] Then these local quaternions are rotated into the world coordinate system using the previous orientation, $Q[t-1]$, of the IMU:

$$Q_i = (\cos(s_{\theta}[t]/2), Q[t-1](1, 0, 0)^T Q[t-1]^{-1} \sin(s_{\theta}[t]/2)) \quad (23)$$

$$Q_j = (\cos(s_{\phi}[t]/2), Q[t-1](0, 1, 0)^T Q[t-1]^{-1} \sin(s_{\phi}[t]/2)) \quad (24)$$

$$Q_k = (\cos(s_{\psi}[t]/2), Q[t-1](0, 0, 1)^T Q[t-1]^{-1} \sin(s_{\psi}[t]/2)) \quad (25)$$

[0120] and combine them together into one quaternion representing the next orientation.

$$Q_{\text{gyro}}[t] = Q_i Q_j Q_k \quad (26)$$

[0121] (6) Method for Comparing Motion Signatures

[0122] The method 24 (see **FIG. 2**) presented in this Section solves the problem of how to compare unsynchronized motion signatures that are collected using IMU's with unknown orientations. For example, a typical application would require matching a realtime signal to a database of motion signatures. However, the orientation of the IMU might be arbitrarily different from the orientation of the IMU used to generate the motion signatures in the database. Furthermore, the proper way to temporally align a test signal to the motion signatures in the database must be determined.

[0123] A rotation-invariant Viterbi algorithm (RIVIT) according to the present embodiment is designed for the purpose of finding the optimal rotation and temporal correspondence between two motion signatures. Suppose the two sequences under consideration are represented as matrices whose columns are the stacked values of the acceleration and angular components,

$$X = \begin{pmatrix} U_X \\ V_X \end{pmatrix} = \begin{pmatrix} u_X[0] & \dots & u_X[T_X - 1] \\ v_X[0] & \dots & v_X[T_X - 1] \end{pmatrix} \quad (27)$$

$$Y = \begin{pmatrix} U_Y \\ V_Y \end{pmatrix} = \begin{pmatrix} u_Y[0] & \dots & u_Y[T_Y - 1] \\ v_Y[0] & \dots & v_Y[T_Y - 1] \end{pmatrix}$$

[0124] The lengths of sequences X and Y are T_X and T_Y , respectively.

[0125] **FIG. 8(a)** and **FIG. 8(b)** show examples of two multi-dimensional sequences having possibly different lengths and an arbitrary rotation between them. This particular examples are sequences of motion signatures of a sitting action, but the IMU was oriented differently in each case. Further, the sitting motion in **FIG. 8(a)** was faster than that of **FIG. 8(b)**.

[0126] In order to calculate the similarity between these two sequences to within some arbitrary rotation and some arbitrary time-warping, the problem may be stated as the minimization of the following measure:

$$D_{R,C}(X,Y) = \text{tr}((Y-RXC)^T(Y-RXC)) \quad (28)$$

[0127] over its free parameters which are a rotation matrix, R, and a correspondence matrix, C, where the symbol "tr" means a transverse matrix. A correspondence matrix is a

special kind of projection matrix that maps the time dimension (columns) of X onto the time dimension (columns) of Y. Suppose we had a function, $c(t)$, that assigns each integer in $\{0, 1, \dots, T_X-1\}$ to an integer in $\{0, 1, \dots, T_Y-1\}$. Its correspondence matrix can be constructed as follows:

$$C_{ij} = \begin{cases} 0 & c(i) \neq j \\ 1 & c(i) = j \end{cases} \quad (29)$$

[0128] While minimizing $D_{R,C}(X, Y)$, it is preferable to constrain C to matrices where the corresponding $c(t)$ is non-decreasing. **FIG. 8(c)** gives an example of a possible mapping. The optimal C map determines how X should be stretched and compressed (possible deleting some sections and repeating others) to line up with Y. The optimal rotation matrix, R, determines how all the vectors in X should be rotated so that they match the vectors in Y.

[0129] The method used in here to minimize the expression (28) is called the coordinate descent method (Non-patent Document 1). It is a simple iterative method that is easily adapted for parallel computation. The idea is to minimize the error function along one single axis in each iteration. This greatly simplifies the minimization problem, as it can be seen, and when it is cycled through each axis, successively minimizing each coordinate, it is known that the Equation will converge to a local minimum. If there were no constraints on R and C then since $D_{R,C}(X, Y)$ is convex this local minimum would also be the global minimum. However, there is constraints in the present embodiment so only localness is guaranteed. In the next Sections, it is described how to minimize the error function over each axis, R and C.

[0130] (6.1) Minimizing $D_{R,C}(X, Y)$ over R

[0131] There is a simple method for minimizing $D_{R,C}(X, Y)$ over just R. Lets define:

$$R^*(C) \equiv \underset{R}{\text{argmin}} D_{R,C}(X, Y) \quad (30)$$

[0132] The solution to this minimization problem is well-known as polar decomposition (see Non-patent Document 6 for a detailed derivation) and the solution is:

$$R^*(C) = VU^T \quad (31)$$

[0133] Matrices V and U are obtained from the singular value decomposition (SVD) of $XC Y^T = UDV^T$. It should be noted that the size of $XC Y^T$ is 6-by-6 so that the cost of the SVD is independent of the length of the sequences.

[0134] The inventor of the present application have found that it is beneficial to restrict R so that the same rotation is applied to the gyro and accelerometer components. This means that the structure of R is desirable to be a block diagonal like this:

$$R = \begin{pmatrix} Q & 0 \\ 0 & Q \end{pmatrix} \quad (32)$$

[0135] where Q is 3-by-3 rotation matrix. This constraint can be easily satisfied by calculating $Q=VU'$ but this time V and U are obtained from the SVD of $U_X C U_Y' + V_X C V_Y' = U D V'$.

[0136] (6.2) Minimizing $D_{R,C}(X, Y)$ over C

[0137] There is also a method for minimizing $D_{R,C}(X, Y)$ over just C using dynamic time warping (DTW). As before, lets define:

$$C^*(R) = \arg \min_C D_{R,C}(X, Y) \quad (33)$$

[0138] Instead of finding C directly as a matrix, the underlying mapping function, $c(t)$ is calculated. Let the vectors of Y sequence represent the states in a Hidden Markov Model (HMM) with left-right topology. Sequence X is matched against the HMM using Viterbi algorithm to provide the optimal path sequence. Scoring of the i-th HMM-state (which represents the i-th vector of Y) on the j-th vector of X is calculated as the squared Euclidean distance:

$$\log P(s_i=i|Y_i, X_j) = -(RX_j - Y_i)(RX_j - Y_i) \quad (34)$$

[0139] This is equivalent to using unit variance Gaussians centered on the vectors of X to represent the output distributions of each HMM-state, $s_i \in \{0, 1, \dots, TY-1\}$. The transition probabilities are:

$$\log P(s_i = j | s_{i-1} = j) = \begin{cases} \log \epsilon & j < i \\ -\log 2 & j = i \\ -\frac{K(i-j-1)^2}{r_{IMU}} & j > i \end{cases} \quad (35)$$

[0140] where K is an experimentally determined constant (approx. (100/128)Hz in our experiments) and r is some very small number. Alternatively, the transition probabilities may be determined by using training data.

[0141] Using the HMM with parameters specified above as representing Y and an observation sequence as X, the standard Viterbi algorithm is used to find the optimal (maximum likelihood) state path, $\{s_t^*\}$ through the HMM. This state path assigns each vector in X to a vector in Y, according to $c(t)=s_t^*$. The correspondence matrix can then be calculated directly from $c(t)$ as stated above. It should be noted that, in an actual implementation, the matrix product XC is much more efficiently calculated by using $c(t)$ directly to construct XC from X.

[0142] (6.3) Determining Initial Conditions

[0143] The coordinate descent method, like all descent methods, requires an initial value for at least one of the coordinates. The quality of the local minimum found by coordinate descent is highly sensitive to this initial value.

The optimization of $D_{R,C}(X, Y)$ is started with an initial guess for C. The initial value for C is calculated using the same procedure as in Section (6.2) but equation (34) is replaced with,

$$\log P(s_i=i|Y_i, X_j) = -(X_j - Y_i)(X_j - Y_i)^2 \quad (36)$$

[0144] so that there is no dependence on the rotation matrix, R. This way no rotation matrix is required to find an approximate mapping, $c(t)$, between X and Y.

[0145] (6.4) Calculating the Distortion Between X and Y

[0146] Below, the procedure, called DISTORTION, is presented for minimizing the error function and thereby calculating the distortion measure between two motion signatures. Let i represent the iteration, and C^0 represent the initial guess for C as calculated in Section (6.3) and ϵ is a stopping criterion used to detect convergence:

```

DISTORTION(X, Y, C0, ε) returns Dfinal
1. Initialize R0 ← R*(C0), δ ← 1, i ← 1
2. While δ > ε:
3.   Ci ← C*(Ri-1)
4.   Ri ← R*(Ci)
5.   Di ← DRi, Ci(X, Y)
6.   δ ← |Di - Di-1|
7.   i ← i + 1
8. Dfinal ← Di
    
```

[0147] $R^*(C)$ and $C^*(R)$ are the procedures given in Sections (6.1) and (6.2), respectively. In summary, the DISTORTION procedure calculates a local minimum of $D_{R,C}(X, Y)$ over R and C using the coordinate descent method outlined above. The value returned, D^{final} , is a measure of the dissimilarity between X and Y that is invariant to any rotation or time-warping.

[0148] (6.5) Device Configuration

[0149] The method described in this Section for comparing motion signatures may be realized by the computational unit 20, which may include a CPU and memory for performing processing steps of the method. More specifically, for example, the computational unit 20 may include a buffer 22, a Rotation Invariant Viterbi (RIVIT) processing section 24, a Hidden Markov Model (HMM) storage 26 and an output section 28.

[0150] The buffer 22 receives and stores sampled sensor outputs to construct a observation sequence X indicative of motion signature corresponding to a current activity of the subject. The observation sequence X is sent to the RIVIT processing section 24 for finding the optimal rotation and temporal correspondence with a sequence Y, one of sequences represented by the HMM stored in the HMM storage 26. In the RIVIT processing section 24, the optimal rotation may be determined first, and then the optimal temporal correspondence may be determined by matching the sequence X against the HMM using Viterbi to provide the optimal path sequence.

[0151] The output section 28 receives the matching result, and outputs as an activity 10 recognition result a signal indicating the current activity of the User. In the present embodiment, it is assumed that the activities may be labeled or classified in advance and mapped against sequences or

states of the HMM. The mapping table may be stored in the output section 28 or the sequences or states of the HMM may be labeled.

[0152] (7) Method for Determining the Location of the Sensor on the Body

[0153] In this Section the method 22 for estimating the location of the IMU with respect to the User's body. Any given movement by the User will generate differing measurements depending on where the IMU is on the User's body, such as in a pocket, being held in 20 the left hand, on a necklace, etc. These differences in measurement of the same action can be used to deduce the location of the IMU. Of course, for this purpose some types of movement are better than others. For example, a motion that only involves the left hand can not be used to distinguish between one situation where the IMU is in the User's left pocket and the situation where the IMU is in the User's right pocket. However, a motion like walking produces unique measurements all over the User's body, except for left-right symmetry.

[0154] In this Section, the action or movement used to determine IMU location is referred to as the reference motion. The overall processing pipeline, assuming an arbitrary reference motion X, is depicted in FIG. 5.

[0155] A position-invariant detector of the reference motion is used to gate the output of the IMU (50), and this way a sequence of test signatures containing the reference motion arrive at the next module 24. For example, if the reference motion is walking then the method 21 for gait frequency measurement (Section (4)) can be used as the position-invariant detector.

[0156] In the module 24, the test signatures are matched against an exhaustive database 54 of reference motion signatures over locations. The best match indicates the estimated location of the IMU. For generic matching of motion signatures the method in Section (6) may be used. It is also possible to custom-design a matching algorithm for a specific reference motion, especially if it is a periodic motion such as walking.

[0157] According to the present embodiment, in its most basic instantiation, the Device is continually outputting what it calculates as the motion that the User has just performed or is performing. What it specifically outputs is determined by the database of motions that are available to the Device for matching. This motion database can be either static or dynamic depending on application needs.

[0158] It is also possible to use the Device in a configuration where it is not attached to any human user but rather to a machine or any moving object. Examples of this include vehicles (such as bicycles, cars and boats), factory equipment (such as robotic arms and CNC machinery), plants and animals (such as dogs and cats), cargo (such as a crate of eggs or a drum of volatile liquid).

[0159] The classification results can be used directly to drive an application that needs to configure itself based on what the User has been doing or is doing at the moment. These type of applications would be called context-aware and their context-awareness is enabled by, in this case, the Device.

[0160] The present invention contains subject matter related to Japanese Patent Application JP 2005-169507 filed in the Japanese Patent Office on Jun. 9, 2005, the entire contents of which being incorporated herein by reference.

[0161] It should be understood by those skilled in the art that various modifications, combinations, sub-combinations and alterations may occur depending on design requirements and other factors insofar as they are within the scope of the appended claims or the equivalents thereof.

What is claimed is:

1. An activity recognition apparatus for detecting an activity of a subject, the apparatus comprising:

a sensor unit including a plurality of linear motion sensors configured to detect linear motions and a plurality of rotational motion sensors, the linear motions being orthogonal to each other, the rotational motions being orthogonal to each other; and

a computational unit configured to receive and process signals from the sensors included in the sensor unit so as to detect an activity of the subject;

wherein the sensor unit is directly or indirectly supported by the subject with an arbitrary orientation with respect to the subject; and

wherein the computational unit performs a calculation that uses the signals from both linear motion sensors and rotational motion sensors to determine the activity of the subject independent of the orientation of the sensor unit.

2. The activity recognition apparatus according to claim 1, wherein:

the plurality of linear motion sensors are three accelerometers for measuring accelerations in three orthogonal axes;

the plurality of rotational motion sensors are three gyro sensors for measuring angular velocities around the three orthogonal axes; and

the computational unit uses six outputs from the accelerometers and the gyro sensors to determine a current activity of the subject.

3. The activity recognition apparatus according to claim 1 or 2, wherein:

the computational unit includes

a buffer configured to store an observation sequence formed of a plurality set of the sensor outputs for a predetermined period of time;

a storage configured to store a plurality of reference sequences, the reference sequences corresponding to motion signatures of different activities; and

a matching processor configured to match the observation sequence against the reference sequences to find a best matched reference sequence;

wherein the matching processor finds an optimal rotation and temporal correspondence between the observation sequence and one of the reference sequences so as to obtain the best matched reference sequence.

4. The activity recognition apparatus according to claim 3, wherein:

the reference sequence includes a plurality of vectors representing states in Hidden Markov Model; and

the matching between the observation sequence and the reference sequence is performed using Viterbi algorithm so as to obtain an optimal state path.

5. The activity recognition apparatus according to claim 1 or 2, wherein:

the computational unit further detects a rate of the subject's gait by calculating an autocorrelation of the signals from the linear motion sensors and an autocorrelation of the signals from the rotation motion sensors, and summing these autocorrelations to obtain a summary autocorrelation function; and

the summary autocorrelation function is used to determine the rate of the subject's gait.

6. The activity recognition apparatus according to claim 1 or 2, wherein:

the computational unit further detects a rotation of the sensor unit by calculating a first rotation by integrating the signals from the rotational motion sensors, calculating a second rotation by using a gravitation model and the signals from the linear motion sensors, and performing weighted interpolation of the first and second rotation;

the weight for the first rotation is larger if variations in the signals from the rotational motion sensors are smaller; and

the weight for the second rotation is larger if variations in the signals from the linear motion sensors are smaller.

7. The activity recognition apparatus according to claim 1 or 2, further comprising:

a storage configure to store reference motion signatures and corresponding locations for a reference activity; wherein

the computational unit further detects a location of the sensor unit on the subject by determining if a current

activity represented with the signals from the sensors is the reference activity, and, if the current activity is the reference activity, matching a motion signature corresponding to the current activity against the reference motion signature so as to determine a location of the sensor unit, which corresponds to a best matched reference motion signature.

8. An activity recognition method of determining an activity of a subject, the method comprising:

sampling accelerations in three orthogonal axes and angular velocities around the three orthogonal axes; and

determining an activity of the subject using all the sampled accelerations and angular velocities;

wherein the accelerations and angular velocities are detected by sensors integrated into a single unit, which is directly or indirectly supported by the subject with an arbitrary orientation with respect to the subject; and

wherein the determination of activity is performed by using both the accelerations and angular velocities to determine the activity independent of the orientation of the sensor unit.

9. A recording medium that stores a program that causes a computer to perform activity recognition by determining an activity of a subject, the program comprising:

receiving data of accelerations in three orthogonal axes and angular velocities around the three orthogonal axes; and

determining an activity of the subject using all the received data of accelerations and angular velocities;

wherein the accelerations and angular velocities are detected by sensors integrated into a single unit, which is directly or indirectly supported by the subject with an arbitrary orientation with respect to the subject; and

wherein the determination of activity is performed by using both the accelerations and angular velocities to determine the activity independent of the orientation of the sensor unit.

* * * * *