

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
07 June 2018 (07.06.2018)

(10) International Publication Number
WO 2018/102190 A1

(51) International Patent Classification:
G01C 3/00 (2006.01) **H04N 7/18** (2006.01)
G06K 9/62 (2006.01)

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(21) International Application Number:
PCT/US2017/062721

(22) International Filing Date:
21 November 2017 (21.11.2017)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/427,573 29 November 2016 (29.11.2016) US


(71) Applicant: **BLACKMORE SENSORS AND ANALYTICS INC.** [US/US]; 3991 Valley Commons Dr, Bozeman, Montana 59718 (US).

(72) Inventors: **CROUCH, Stephen C.**; 127 E Alderson St., Apt A, Bozeman, Montana 59715 (US). **REIBEL, Randy R.**; 3570 W. Broadwater, Bozeman, Montana 59718 (US). **KAYLOR, Brant**; 32 Erik St., Merrimack, New Hampshire 03054 (US).

(74) Agent: **MOLINELLI, Eugene J.** et al.; P.O. Box 231, Manassas, Virginia 20108 (US).

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

(54) Title: METHOD AND SYSTEM FOR CLASSIFICATION OF AN OBJECT IN A POINT CLOUD DATA SET



(57) Abstract: A method for classifying an object in a point cloud includes computing first and second classification statistics for one or more points in the point cloud. Closest matches are determined between the first and second classification statistics and a respective one of a set of first and second classification statistics corresponding to a set of N classes of a respective first and second classifier, to estimate the object is in a respective first and second class. If the first class does not correspond to the second class, a closest fit is performed between the point cloud and model point clouds for only the first and second classes of a third classifier. The object is assigned to the first or second class, based on the closest fit within near real time of receiving the 3D point cloud. A device is operated based on the assigned object class.

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *with international search report (Art. 21(3))*

METHOD AND SYSTEM FOR CLASSIFICATION OF AN OBJECT IN A POINT CLOUD DATA SET

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims benefit of Provisional Appln. 62/427,573, filed November 29, 2016, the entire contents of which are hereby incorporated by reference as if fully set forth herein, under 35 U.S.C. §119(e).

STATEMENT OF GOVERNMENTAL INTEREST

[0002] This invention was made with government support under contract N00014-16-C-1026 awarded by the Department of Navy. The government has certain rights in the invention.

BACKGROUND

[0003] Optical detection of range, often referenced by a mnemonic, LIDAR, for light detection and ranging, is used for a variety of applications, from altimetry, to imaging, to collision avoidance. LIDAR provides finer scale range resolution with smaller beam sizes than conventional microwave ranging systems, such as radio-wave detection and ranging (RADAR). Optical detection of range can be accomplished with several different techniques, including direct ranging based on round trip travel time of an optical pulse to a target, and chirped detection based on a frequency difference between a transmitted chirped optical signal and a returned signal scattered from a target, and phase encoded detection based on a sequence of single frequency phase changes that are distinguishable from natural signals.

[0004] To achieve acceptable range accuracy and detection sensitivity, direct long range LIDAR systems use short pulse lasers with low pulse repetition rate and extremely high pulse peak power. The high pulse power can lead to rapid degradation of optical components. Chirped LIDAR systems use long optical pulses with relatively low peak optical power. In this configuration, the range accuracy depends on the chirp bandwidth rather than the pulse duration, and therefore excellent range accuracy can still be obtained.

[0005] Useful optical chirp bandwidths have been achieved using wideband radio frequency (RF) electrical signals to modulate an optical carrier. Recent advances in chirped LIDAR include using the same modulated optical carrier as a reference signal that is combined with the returned signal at an optical detector to produce in the resulting electrical signal a

relatively low beat frequency that is proportional to the difference in frequencies between the references and returned optical signals. This kind of beat frequency detection of frequency differences at a detector is called heterodyne detection. It has several advantages known in the art, such as the advantage of using RF components of ready and inexpensive availability. Recent work described in U.S. Patent Number 7,742,152, shows a novel simpler arrangement of optical components that uses, as the reference optical signal, an optical signal split from the transmitted optical signal. This arrangement is called homodyne detection in that patent.

[0006] LIDAR detection with phase encoded microwave signals modulated onto an optical carrier have been used as well. This technique relies on correlating a sequence of phases (or phase changes) of a particular frequency in a return signal with that in the transmitted signal. A time delay associated with a peak in correlation is related to range by the speed of light in the medium. Advantages of this technique include the need for fewer components, and the use of mass produced hardware components developed for phase encoded microwave and optical communications.

[0007] The data returned by these LIDAR systems is often represented as a point cloud. A point cloud is a set of data points in some coordinate system. In a three dimensional coordinate system, these points are usually defined by X, Y and Z coordinates, and often are intended to represent the external surface of an object. 3D point clouds can be generated by 3D scanners, such as LIDAR systems including chirped LIDAR and phase coded LIDAR, among other types of scanners.

SUMMARY

[0008] The current inventors have recognized circumstances and applications in which automatic classification of objects represented by 3D point clouds is challenging in real time, particularly for objects located at long range. Techniques are provided for such automatic classification of objects.

[0009] In a first set of embodiments, a method implemented on a processor includes obtaining a 3D point cloud representing an external surface of an object. The method further includes extracting a surface normal at a point on the 3D point cloud to define a translational and rotational invariant coordinate system at the point and extracting values of one or more feature variables in the coordinate system to define one or more points in the 3D point cloud. The method further includes computing a first classification statistic for the one or more

points in the 3D point cloud based on values of the one or more feature variables and computing a second classification statistic for the one or more points in the 3D point cloud based on values of the one or more feature variables. The method further includes determining a closest match between the first classification statistic and a first one of a set of N first classification statistics corresponding to a set of N classes of a first classifier to estimate that the object belongs in a first class. The method further includes determining a closest match between the second classification statistic and a second one of a set of second classification statistics corresponding to the set of N classes of a second classifier to estimate that the object belongs in a second class. If the first class does not correspond to the second class, a third classification statistic is computed for the one or more points in the 3D point cloud based on a closest fit between the 3D point cloud and model point clouds for only the first class and second class of a third classifier. The object is assigned to the first class or second class, based on the closest fit of the third classifier within near real time from receiving the 3D point cloud. The method further includes operating a device based on the one class that the object is assigned to.

[0010] In some embodiments of the first set, the first classification statistic is a spin image, the second classification statistic is a covariance matrix, and the third classification statistic is an iterative closest point (ICP).

[0011] In a second set of embodiments, an apparatus includes a laser source configured to provide an optical signal. The apparatus includes a splitter configured to receive the signal and produce a transmitted signal and a reference signal. The apparatus also includes an optical coupler configured to direct the transmitted signal outside the apparatus and to receive any return signal backscattered from any object illuminated by the transmitted signal. The apparatus also includes an optical detector disposed to receive the reference signal and the return signal. In addition, the apparatus still further includes a processor configured to perform the steps of receiving an electrical signal from the optical detector. The processor is further configured to perform one or more steps of the above method.

[0012] In other embodiments, a system or apparatus or computer-readable medium is configured to perform one or more steps of the above methods.

[0013] Still other aspects, features, and advantages are readily apparent from the following detailed description, simply by illustrating a number of particular embodiments and implementations, including the best mode contemplated for carrying out the invention. Other

embodiments are also capable of other and different features and advantages, and its several details can be modified in various obvious respects, all without departing from the spirit and scope of the invention. Accordingly, the drawings and description are to be regarded as illustrative in nature, and not as restrictive.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] Embodiments are illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings in which like reference numerals refer to similar elements and in which:

[0015] FIG. 1A is a set of graphs that illustrates an example optical chirp measurement of range, according to an embodiment;

[0016] FIG. 1B is a graph that illustrates an example measurement of a beat frequency resulting from de-chirping, which indicates range, according to an embodiment;

[0017] FIG. 2 is a block diagram that illustrates example components of a high resolution LIDAR system, according to an embodiment;

[0018] FIG. 3A is a block diagram that illustrates example components of a heterodyne chirped LIDAR system, according to an embodiment;

[0019] FIG. 3B is a block diagram that illustrates example components of a homodyne chirped LIDAR system, according to an embodiment;

[0020] FIG. 4 illustrates an example of a 3D point cloud of an object, according to an embodiment;

[0021] FIG. 5A is a graph that illustrates an example of a k-d tree composition for a point set, according to an embodiment;

[0022] FIG. 5B illustrates an example of a k-d tree based on the k-d tree composition of FIG. 5A, according to an embodiment;

[0023] FIG. 6A illustrates an example of a 3D point cloud of an object, according to an embodiment;

[0024] FIG. 6B illustrates a segment of the 3D point cloud of FIG. 6A, according to an embodiment;

[0025] FIG. 6C is a block diagram that illustrates an example of feature variables to define the points of a 3D point cloud within a rotational and translational invariant coordinate system, according to an embodiment;

[0026] FIG. 6D is a spin image that illustrates an example of a histogram of the number of points of the 3D point cloud over a range of the feature variables of FIG. 6C, according to an embodiment;

[0027] FIG. 6E is a graph that illustrates an example of a plurality of clusters and cluster centers associated with a respective plurality of classes, according to an embodiment;

[0028] FIG. 7 is a flow chart that illustrates an example method for classifying an object defined by a 3D point cloud, according to an embodiment;

[0029] FIG. 8A is a photograph that illustrates an example of a top perspective view of an experimental setup of the system of FIG. 2, according to embodiments;

[0030] FIG. 8B is a block diagram that illustrates an example of objects used in the experimental setup depicted in FIG. 8A, according to an embodiment;

[0031] FIG. 9A is a set of point clouds for a set of objects, according to an embodiment;

[0032] FIG. 9B is a block diagram that illustrates an example of the set of objects used to obtain the set of point clouds of FIG. 9A, according to an embodiment;

[0033] FIG. 10A is a block diagram that illustrates an example of an occluded model of an object based on the 3D point cloud of the object in FIG. 4, according to an embodiment;

[0034] FIG. 10B is a block diagram that illustrates an example of a 2D point cloud of an object, according to an embodiment;

[0035] FIGS. 10C-10E is a block diagram that illustrates an example of a best fit between the 2D point cloud of FIG. 10B and each occluded model of FIG. 10A, according to an embodiment;

[0036] FIG. 11 is a block diagram that illustrates a computer system upon which an embodiment of the invention may be implemented; and

[0037] FIG. 12 illustrates a chip set upon which an embodiment of the invention may be implemented.

DETAILED DESCRIPTION

[0038] A method and apparatus and system and computer-readable medium are described for classifying an object in a 3D point cloud. In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

[0039] Notwithstanding that the numerical ranges and parameters setting forth the broad scope are approximations, the numerical values set forth in specific non-limiting examples are reported as precisely as possible. Any numerical value, however, inherently contains certain errors necessarily resulting from the standard deviation found in their respective testing measurements at the time of this writing. Furthermore, unless otherwise clear from the context, a numerical value presented herein has an implied precision given by the least significant digit. Thus a value 1.1 implies a value from 1.05 to 1.15. The term "about" is used to indicate a broader range centered on the given value, and unless otherwise clear from the context implies a broader range around the least significant digit, such as "about 1.1" implies a range from 1.0 to 1.2. If the least significant digit is unclear, then the term "about" implies a factor of two, e.g., "about X" implies a value in the range from 0.5X to 2X, for example, about 100 implies a value in a range from 50 to 200. Moreover, all ranges disclosed herein are to be understood to encompass any and all sub-ranges subsumed therein. For example, a range of "less than 10" can include any and all sub-ranges between (and including) the minimum value of zero and the maximum value of 10, that is, any and all sub-ranges having a minimum value of equal to or greater than zero and a maximum value of equal to or less than 10, e.g., 1 to 4.

[0040] Some embodiments of the invention are described below in the context of classifying objects in 3D point clouds generated by a LIDAR system with a linear frequency modulated optical signal. However, such 3D point clouds need not be generated with LIDAR systems featuring linear chirps and can instead be generated by LIDAR systems where the transmitted signal is modulated in amplitude, frequency or phase or some combination. Embodiments are described in the context of classifying an object within 3D point clouds generated by a system including a single optical beam and its return on a single detector or pair of detectors,

which can then be scanned using any known scanning means, such as linear stepping or rotating optical components or with arrays of transmitters and detectors or pairs of detectors. Other embodiments involve classifying objects in 3D point clouds that are not generated by any specific system but are instead generated by a system that is different from the illustrated embodiments. In still other embodiments, the 3D point clouds could be generated with other sensors including any type of 3D scanning LIDAR system (e.g. Velodyne®, Riegl®, Leica®, Faro®) or any “flash” LIDAR system such as those employing LIDAR or Geiger-mode Avalanche Photodiode (APD) array for depth-imaging. These systems possess sufficient down-range and cross-range resolution to collect reasonable data for generating 3D point clouds for use with the embodiments discussed herein.

1. Chirped Detection Overview

[0041] FIG. 1A is a set of graphs 110, 120, 130, 140 that illustrates an example optical chirp measurement of range, according to an embodiment. The horizontal axis 112 is the same for all four graphs and indicates time in arbitrary units, on the order of milliseconds (ms, 1 ms = 10^{-3} seconds). Graph 110 indicates the power of a beam of light used as a transmitted optical signal. The vertical axis 114 in graph 110 indicates power of the transmitted signal in arbitrary units. Trace 116 indicates that the power is on for a limited pulse duration, τ starting at time 0. Graph 120 indicates the frequency of the transmitted signal. The vertical axis 124 indicates the frequency transmitted in arbitrary units. The trace 126 indicates that the frequency of the pulse increases from f_1 to f_2 over the duration τ of the pulse, and thus has a bandwidth $B = f_2 - f_1$. The frequency rate of change is $(f_2 - f_1)/\tau$

[0042] The returned signal is depicted in graph 130 which has a horizontal axis 112 that indicates time and a vertical axis 124 that indicates frequency as in graph 120. The chirp 126 of graph 120 is also plotted as a dotted line on graph 130. A first returned signal is given by trace 136a, which is just the transmitted reference signal diminished in intensity (not shown) and delayed by Δt . When the returned signal is received from an external object after covering a distance of $2R$, where R is the range to the target, the returned signal start at the delayed time Δt is given by $2R/c$, were c is the speed of light in the medium (approximately 3×10^8 meters per second, m/s). Over this time, the frequency has changed by an amount that

depends on the range, called f_R , and given by the frequency rate of change multiplied by the delay time. This is given by Equation 1a.

$$f_R = (f_2 - f_1)/\tau * 2R/c = 2BR/c\tau \quad (1a)$$

The value of f_R is measured by the frequency difference between the transmitted signal 126 and returned signal 136a in a time domain mixing operation referred to as de-chirping. So the range R is given by Equation 1b.

$$R = f_R c \tau / 2B \quad (1b)$$

Of course, if the returned signal arrives after the pulse is completely transmitted, that is, if $2R/c$ is greater than τ , then Equations 1a and 1b are not valid. In this case, the reference signal is delayed a known or fixed amount to ensure the returned signal overlaps the reference signal. The fixed or known delay time of the reference signal is multiplied by the speed of light, c, to give an additional range that is added to range computed from Equation 1b. While the absolute range may be off due to uncertainty of the speed of light in the medium, this is a near-constant error and the relative ranges based on the frequency difference are still very precise.

[0043] In some circumstances, a spot illuminated by the transmitted light beam encounters two or more different scatterers at different ranges, such a front and a back of a semitransparent object, or the closer and farther portions of an object at varying distances from the LIDAR, or two separate objects within the illuminated spot. In such circumstances, a second diminished intensity and differently delayed signal will also be received, indicated in graph 130 by trace 136b. This will have a different measured value of f_R that gives a different range using Equation 1b. In some circumstances, multiple returned signals are received. In some circumstances, the transmitted beam encounters multiple portions of an object and multiple returned signals received from the multiple portions of the object are used to determine a respective range to each of the multiple portions of the object. In these circumstances, the respective range to each of the multiple portions of the object is used to generate a point cloud of the object. Some embodiments of the present invention are provided to classify the object according to the generated point cloud.

[0044] Graph 140 depicts the difference frequency f_R between a first returned signal 136a and the reference chirp 126. The horizontal axis 112 indicates time as in all the other aligned graphs in FIG. 1A, and the vertical axis 134 indicates frequency difference on a much

expanded scale. Trace 146 depicts the constant frequency f_R measured during the transmitted chirp, which indicates a particular range as given by Equation 1b. The second returned signal 136b, if present, would give rise to a different, larger value of f_R (not shown) during de-chirping; and, as a consequence yield a larger range using Equation 1b.

[0045] A common method for de-chirping is to direct both the reference optical signal and the returned optical signal to the same optical detector. The electrical output of the detector is dominated by a beat frequency that is equal to, or otherwise depends on, the difference in the frequencies and phases of the two signals converging on the detector. A Fourier transform of this electrical output signal will yield a peak at the beat frequency. This beat frequency is in the radio frequency (RF) range of Megahertz (MHz, 1 MHz = 10^6 Hertz = 10^6 cycles per second) rather than in the optical frequency range of Terahertz (THz, 1 THz = 10^{12} Hertz). Such signals are readily processed by common and inexpensive RF components, such as a Fast Fourier Transform (FFT) algorithm running on a microprocessor or a specially built FFT or other digital signal processing (DSP) integrated circuit. In other embodiments, the return signal is mixed with a continuous wave (CW) tone acting as the local oscillator (versus a chirp as the local oscillator). This leads to the detected signal which itself is a chirp (or whatever waveform was transmitted). In this case the detected signal would undergo matched filtering in the digital domain as described in Kachelmyer 1990. The disadvantage is that the digitizer bandwidth requirement is generally higher. The positive aspects of coherent detection are otherwise retained.

[0046] FIG. 1B is a graph that illustrates an example measurement of a beat frequency resulting from de-chirping, which indicates range, according to an embodiment. The horizontal axis 152 indicates frequency in Megahertz; and the vertical axis indicates returned signal power density I_R relative to transmitted power density I_T in decibels (dB, Power in dB = $20 \log(I_R/I_T)$). Trace 156 is the Fourier transform of the electrical signal output by the optical detector, such as produced by a FFT circuit and is based on data published by Adany *et al.*, 2009. The horizontal location of the peak gives f_R that indicates the range, using Equation 1b. In addition, other characteristics of the peak can be used to describe the returned signal. For example, the power value at the peak is characterized by the maximum value of trace 156, or, more usually, by the difference 157 (about 19 dB in FIG. 1B) between the peak value (about -31 dB in FIG. 1B) and a noise floor (about -50dB in FIG. 1B) at the shoulders

of the peak; and, the width of the peak is characterized by the frequency width 158 (about 0.08 MHz in FIG. 1B) at half maximum (FWHM). If there are multiple discernable returns, there will be multiple peaks in the FFT of the electrical output of the optical detector, likely with multiple different power levels and widths. Any method may be used to automatically identify peaks in traces, and characterize those peaks by location, height and width. For example, in some embodiments, FFTW or Peak detection by MATLAB – Signal Processing Toolbox is used, available from MATLAB™ of MATHWORKS™ of Natick, Massachusetts. One can also use custom implementations that rely on FFTW in CUDA and custom peak detection in CUDA™ available from NVIDIA™ of Santa Clara, California. Custom implementation have been programmed on field programmable gate arrays (FPGAs). Commonly used algorithm is to threshold the range profile and run a center of mass algorithm, peak fitting algorithm (3-point Gaussian fit), or nonlinear fit of the peak for some function (such as a Gaussian) to determine the location of the peak more precisely. A moving object can introduce a Doppler frequency shift that causes an offset in the computed range. In some embodiments, Doppler compensation is used. Any known Doppler compensation method and hardware may be used.

[0047] A new independent measurement is made at a different angle, or translated position of a moving LIDAR system, using a different pulse after an interlude of t_i , so that the pulse rate (PR) is given by the expression $1/(\tau+t_i)$. A frame is a 2 dimensional image of ranges in which each pixel of the image indicates a range to a different portion of an object viewed by the transmitted beam. For a frame assembled from transmitted signals at each of 1000 horizontal vertical angles by 1000 vertical angles, the frame includes 10^6 pixels and the frame rate (FR) is 10^{-6} of the pulse rate, e.g., is $10^{-6}/(\tau+t_i)$.

2. Range Detection Hardware Overview

[0048] In order to depict how the range detection approach is implemented, some generic and specific hardware approaches are described. FIG. 2 is a block diagram that illustrates example components of a high resolution LIDAR system, according to an embodiment. A laser source 212 emits a carrier wave 201 that is amplitude or frequency or phase modulated, or some combination, in the modulator 214 to produce a pulse that has a bandwidth B and a duration τ . A splitter 216 splits the chirp into a transmitted beam 205 with most of the energy of the beam 203 and a reference beam 207 with a much smaller amount of energy that is

nonetheless enough to produce good heterodyne or homodyne interference with the returned light 291 scattered from a target (not shown). In some embodiments, the transmitted beam is scanned over multiple angles to profile any object in its path. The reference beam is delayed in a reference path 220 sufficiently to arrive at the detector array 230 with the scattered light. In some embodiments, the splitter 216 is upstream of the modulator 214, and the reference beam 207 is unmodulated. In various embodiments, from less to more flexible approaches, the reference is caused to arrive with the scattered or reflected field by: 1) putting a mirror in the scene to reflect a portion of the transmit beam back at the detector array so that path lengths are well matched; 2) using a fiber delay to closely match the path length and broadcast the reference beam with optics near the detector array, as suggested in FIG. 2, with or without a path length adjustment to compensate for the phase difference observed or expected for a particular range; or, 3) using a frequency shifting device (acousto-optic modulator) or time delay of a local oscillator waveform modulation to produce a separate modulation to compensate for path length mismatch; or some combination. In some embodiments, the target is close enough and the pulse duration long enough that the returns sufficiently overlap the reference signal without a delay. In various embodiments, multiple portions of the target scatter a respective returned light 291 signal back to the detector array 230 for each scanned beam resulting in a point cloud based on the multiple ranges of the respective multiple portions of the target illuminated by multiple beams and multiple returns.

[0049] The detector array is a single detector or a 1D or 2D array of detectors arranged in a plane roughly perpendicular to returned beams 291 from the target. The phase or amplitude of the interface pattern, or some combination, is recorded by acquisition system 240 for each detector at multiple times during the pulse duration τ . The number of temporal samples per pulse duration affects the down-range extent. The number is often a practical consideration chosen based on pulse repetition rate and available camera frame rate. The frame rate is the sampling bandwidth, often called “digitizer frequency.” Basically, if X number of detector array frames are collected during a pulse with resolution bins of Y range width, then a $X*Y$ range extent can be observed. The acquired data is made available to a processing system 250, such as a computer system described below with reference to FIG. 11, or a chip set described below with reference to FIG. 12. In some embodiments, the acquired data is a point cloud based on the multiple ranges of the respective multiple portions of the target. An object classification statistic module 270 classifies the object illuminated with the beam 205

based on the acquired point cloud, according to the method 700 of FIG. 7. Any known apparatus or system may be used to implement the laser source 212, modulator 214, beam splitter 216, reference path 220, detector array 230, or acquisition system 240. Optical coupling to scan or flood or focus on the target or focus past a pupil plane are not depicted. As used herein, an optical coupler is any component that affects the propagation of light within spatial coordinates to direct light from one component to another component, such as a vacuum, air, glass, crystal, mirror, lens, optical circulator, beam splitter, phase plate, polarizer, and optical fiber, among others, alone or in some combination

[0050] For example, in some chirp embodiments, the laser used was actively linearized with the modulation applied to the current driving the laser. Experiments were also performed with electro-optic modulators providing the modulation. The system is configured to produce a chirp of bandwidth B and duration τ , suitable for the down-range resolution desired, as described in more detail below for various embodiments. For example, in some illustrated embodiments, a value of B about 90 GHz and τ of about 200 milliseconds (ms, $1 \text{ ms} = 10^{-3}$ seconds) were chosen to work within the confines of the relatively low detector array frame rate in the experiments performed. These choices were made to observe a reasonably large range window of about 30 cm, which is often important in shape of an object and identification of the object. This technique will work for chirp bandwidths from 10 MHz to 5 THz. However, for the 3D imaging applications, typical ranges are chirp bandwidths from about 300 MHz to about 20 GHz, chirp durations from about 250 nanoseconds (ns, $1 \text{ ns} = 10^{-9}$ seconds) to about 1 millisecond (ms, $1 \text{ ms} = 10^{-3}$ seconds), ranges to targets from about 0 meters to about 20 km, spot sizes at target from about 3 millimeters (mm, $1 \text{ mm} = 10^{-3}$ meters) to about 1 meter (m), depth resolutions at target from about 7.5 mm to about 0.5 m. In some embodiments, the target has a minimum range, such as 400 meters (m). It is noted that the range window can be made to extend to several kilometers under these conditions and that the Doppler resolution can also be quite high (depending on the duration of the chirp). Although processes, equipment, and data structures are depicted in FIG. 2 as integral blocks in a particular arrangement for purposes of illustration, in other embodiments one or more processes or data structures, or portions thereof, are arranged in a different manner, on the same or different hosts, in one or more databases, or are omitted, or one or more different

processes or data structures are included on the same or different hosts. For example splitter 216 and reference path 220 include zero or more optical couplers.

[0051] FIG. 3A is a block diagram that illustrates example components of a heterodyne chirped LIDAR system 300a, according to an embodiment. This system 300a, modified from that described in U.S. Patent Number 7,742,152, uses electronic de-chirping. Although an object 390 is depicted to illustrate operation of the system 300a, the object 390 is not part of the system 300a. System 300a includes laser 301, modulator 310, telescope as a scanning optical coupler 320, balanced photodetector 330, processing circuitry 340, waveform generator 350 that produces an FM chirp of bandwidth B and duration τ , power splitter 351, de-chirping mixer 360, and acoustic-optic modulator 370. In this system, the source laser 301 output beam is split by beam splitter 302 into two parts; one part is modulated by modulator 310 based on the FM chirp from power splitter 351 and operational amplifier 352a to produce beam 305 that is fed to the telescope.

[0052] The other part of the beam, beam 307a is used to generate a local oscillator (LO) for coherent detection. An acoustic speaker produces an acoustic signal with frequency fm to drive an acousto-optic modulator (AOM) 370 to shift the optical frequency by fm in beam 307b, which serves as an intermediate frequency (IF) for heterodyne detection. Optical coupler 322 directs beam 307b onto one of the balanced photodetector 330.

[0053] A return optical signal 391 is also directed by optical coupler 322 to the other part of the balanced photodetector. The balanced photodiode 330 rejects the direct detection component. The output electrical signal is amplified in operational amplifier 344a and the IF signal is selected by a bandpass filter 341 and detected by a Schottky diode 342 which recovers the baseband waveform. The resulting electrical signal is directed through low pass filter 343 and operational amplifier 344b.

[0054] A de-chirping mixer 360 compares this detected signal with the original chirp waveform output by power splitter 351 and operational amplifier 352b to produce an electrical signal with the beat frequency that depends on the frequency difference between the RF reference waveform and the detected waveform. Another operational amplifier 344c and a FFT process 345 is used to find the beating frequency. Processor 346 is programmed to do data analysis. Coherent detection systems like 300a significantly improve receiver signal to noise ratio (SNR) compared to direct detection of pulse travel time, however, at the cost of greatly increased system complexity. The electrical components from operational amplifier

344a and de-chirping mixer 360 through processor 346 constitute a signal processing component 340.

[0055] According to the illustrated embodiment, the light beam emitted from optical coupler 320 impinges on one or more objects 390 with a finite beam size that illuminates an illuminated portion 392 of the one or more objects. Backscattered light from an illuminated portion is returned through the telescope to be directed by optical coupler 322 onto the optical detector, such as one photodiode of a balanced photodetector 330. In some embodiments, the light beam emitted from optical coupler 320 impinges on multiple portions 392 of an object and backscattered light 391 is returned from each of the multiple portions 392 of the object to determine the respective range to each multiple portion of the object. In these embodiments, the respective range of each multiple portion of the object is used to add a point to a point cloud of one or more objects. After many measurements at different angles or different positions of the LIDAR system, a point cloud emerges. The processor 346 includes an object classification statistic module 380, as described below, to classify the object 390 illuminated with the beam 305 based on the point cloud, according to the method 700 of FIG. 7.

[0056] FIG. 3B is a block diagram that illustrates example components of a homodyne chirped LIDAR system 300b, according to an embodiment. This system 300b, modified from that described in U.S. Patent Number 7,742,152, uses photonic de-chirping and simplifies the RF components. Although an object 390 is depicted to illustrate operation of the system 300b, the object 390 is not part of the system 300b. The system 300b includes waveform generator 350, laser 301, modulator 310, splitter 302 downstream of the modulator 310, telescope used as scanning optical coupler 320, balanced photodetector 330, and processing circuitry 360.

[0057] In this system, both the optical signal and the local oscillator LO are driven by the same waveform generator 350 and amplified in operational amplifier 352. The beam output by the modulator 310 is split by beam splitter 302 to a beam part 305 and a beam part 307c. The beam part 305, with most of the beam energy, e.g., 90% or more, is transmitted through the optical coupler 320 to illuminate the illuminated portion 392 of the object 390. In some embodiments, the beam 305 illuminates multiple portions 392 of the object 390. In these embodiments, multiple return signals 309 are directed to the photodetector 330 from each of the multiple portions 392 of the object 390 and the respective range to each of the multiple portions 392 of the object 390 is determined. The beam part 307c is delayed a desired

amount in delay 308 to produce the reference signal 307d. In some embodiments, there is no delay and delay 308 is omitted. The reference signal 307d and the return signal 309 from the telescope or other optical coupler 320 are directed to the photodetector 330 by optical couplers 322.

[0058] The de-chirping process is accomplished within the balanced photodiode 330 and therefore eliminates the need of de-chirping mixing and the associated RF processing. Because the original chirp optical waveform, which is carried by the LO, beats with its delayed version at the photodiode as indicated, target distance can be directly obtained by a frequency analysis in an FFT component 345 of the photocurrent signal output by operational amplifier 344. The detected target distance, which is the range to the illuminated portion, for each return is added as a point to the point cloud of the one or more objects. In some embodiments, each portion of the target has a minimum range, such as 400 meters (m). Processor 362 is programmed to do data analysis. The processor 362 includes an object classification statistic module 380, as described below, to classify the object illuminated with the beam 305 based on a point cloud, according to the method 700 of FIG. 7. The electrical components from operational amplifier 344 through processor 362 constitute a signal processing component 360. Considering that shot noise is the dominant noise with coherent detection, SNR at the beating frequency is reduced compared to SNR of direct detection and SNR of the system 300a.

3. K-d Tree Overview

[0059] A k-d tree (short for k-dimensional tree) is a space-partitioning data structure for organizing points in a k-dimensional space. K-d trees are a useful data structure for several applications, such as searches involving a multidimensional search key (e.g. range searches and nearest neighbor searches). K-d trees are a special case of binary space partitioning trees. FIG. 5A is a graph that illustrates an example of a k-d tree composition 500 for a point set, according to an embodiment. FIG. 5B illustrates an example of a k-d tree 550 based on the k-d tree composition 500 of FIG. 5A, according to an embodiment. The example embodiments of FIGS. 5A-5B depict a k-d tree where k=2 (e.g. two dimensional x-y space). The k-d tree of FIGS. 5A-5B is based on the point set {(2,3), (5,4), (9,6), (4,7), (8,1), (7,2)}. The graph of FIG. 5A plots each point in the point set. The horizontal axis 502 is the x-axis and the vertical axis 504 is the y-axis.

[0060] The k-d tree 550 is a binary tree with a plurality of nodes 552, 554, 556, including a root node 552 and a leaf node 556. Every non-leaf node (e.g. node 552, 554) can be thought of as implicitly generating a splitting hyperplane that divides the space into two parts, known as half-spaces. Points to the left of this hyperplane are represented by the left subtree of that node and points right of the hyperplane are represented by the right subtree. The hyperplane direction is chosen in the following way: every node in the tree is associated with one of the k-dimensions, with the hyperplane perpendicular to that dimension's axis. In the example embodiment of FIG. 5B, root node 552 is associated with the x-dimension and a point 552a is chosen (e.g. (7,2)) from the point set with a median value in the x-dimension. An x-splitting plane 506 is generated in the k-d tree composition 500 based on this median value in the x-dimension (e.g. x=7). Points from the point set to the left of this plane 506 are represented by the left subtree 553 and points from the point set to the right of this plane 506 are represented by the right subtree 555.

[0061] Node 554 is associated with the y-dimension. A point 554a is chosen (e.g. (5,4)) from the points to one side (left side) of the x-splitting plane 506 (e.g. points in the left subtree 553) with a median value in the y-dimension. A y-splitting plane 508 is generated in the k-d tree composition 500 based on this median value in the y-dimension (e.g. y=4). A point to the one side (bottom side) of this y-splitting plane 508 is represented in the leaf node 556 by point (2,3) while a point to the other side (top side) of this y-splitting plane 508 is represented in the right leaf node 556 by point (4,7). Similarly, a point 554b is chosen (e.g. (9,6)) from the point set in the right subtree 555 with a median value in the y-dimension and a y-splitting plane 510 is generated in the k-d tree composition based on this median value in the y-dimension (e.g. y=6). A point to the left of this y-splitting plane 510 is represented in the leaf node 556 by point (8,1), while there are no points to the right of this y-splitting plane 510 in the point set.

[0062] In some embodiments, the k-d tree 550 can be used to perform a nearest neighbor (NN) search, which aims to find the point in the set that is nearest to a given input point. In an example embodiment, the input point 512 is (2,8). This search can be done efficiently by using the tree properties to quickly eliminate large portions of the search space. In some embodiments, searching for a nearest neighbor in a k-d tree proceeds by starting with the root node 552a. The search moves down the tree 550 recursively, in the same way that it would if the input point were being inserted (i.e. it goes left or right depending on whether the point is

lesser than or greater than the current node in the split dimension). At the root node 552a, the input point (2,8) is less than the point (7,2) at node 552a in the split dimension (e.g. x dimension) and thus the search moves to the left subtree 553. At the node 554a, the input point (2,8) is greater than the point (5,4) at node 554a in the split dimension (e.g. y dimension) and thus the search moves to the right and to point (4,7) at the leaf node 556. Once the search reaches a leaf node 556 point, the search saves the node point as the “nearest neighbor”. In this example, the search saves (4,7) as the “nearest neighbor” to the input point (2,8). FIG. 5A depicts the input point 512 and nearest neighbor 514 corresponding to (4,7) at the leaf node 556. Software is available that can produce K-d trees for K-dimensional datasets. For example, Accord-framework.net is an open source software platform that provides k-d tree nearest neighbor search algorithms in C#.

4. Point Cloud Overview

[0063] A point cloud is a set of data points in some coordinate system. In a three-dimensional coordinate system, these points are usually defined by X, Y, and Z coordinates, and often are intended to represent the external surface of an object. In other embodiments each point in the point cloud is represented in a different coordinate system, such a polar coordinate system that positions each point using a range and a azimuthal angle and an elevation angle from some convenient origin, such as the location of the LIDAR scanner. Point clouds may be created by 3D scanners. These devices measure a large number of points on an object's surface, and often output a point cloud as a data file. The point cloud represents the set of points that the device has measured. FIG. 4 illustrates an example of a 3D point cloud 400 of an object, such as a cup, according to an embodiment. The 3D scanner measures points at different portions of the cup surface, including a portion 402a corresponding to the cup handle and a portion 402b corresponding to the cup rim, for example. In some embodiments, 3D point clouds of objects are created by optical LIDAR systems, such as the systems 200, 300a, 300b of FIGS. 2 and 3A-3B previously discussed above. In these embodiments, a 3D point cloud of the object (not shown) in FIG. 2 or the object 390 in FIGS. 3A-3B is obtained including a number of points on the object surface, based on the respective range calculated to multiple portions of the object. If the object is opaque to the light of the scanner and the object and scanner are immobile, only one side of the object is observed and the point cloud is said to be occluded. If either the object or the

scanner is moved to expose multiple sides of the object to the scanner (e.g., by rotating the object), than a more complete 3D representation of the surface of the object can be obtained.

[0064] As the output of 3D scanning processes, point clouds are used for many purposes, including to create 3D CAD models for manufactured parts, metrology/quality inspection, and a multitude of visualization, animation, rendering and mass customization applications. While point clouds can be directly rendered and inspected, usually point clouds themselves are generally not directly usable in most 3D applications, and therefore are usually converted to polygon mesh or triangle mesh models, NURBS surface models, or CAD models through a process commonly referred to as surface reconstruction. There are many techniques for converting a point cloud to a 3D surface. Some approaches, like Delaunay triangulation, alpha shapes, and ball pivoting, build a network of triangles over the existing vertices of the point cloud, while other approaches convert the point cloud into a volumetric distance field and reconstruct the implicit surface so defined through a marching cubes algorithm.

[0065] One application in which point clouds are directly usable is industrial metrology or inspection using industrial computed tomography. The point cloud of a manufactured part can be aligned to a CAD model (or even another point cloud), and compared to check for differences. These differences can be displayed as color maps that give a visual indicator of the deviation between the manufactured part and the CAD model. Geometric dimensions and tolerances can also be extracted directly from the point cloud.

5. Object Classification statistic Overview

[0066] Object classification seeks to determine from a point cloud of an otherwise unknown object, the class of objects to which the object giving rise to the point cloud probably belongs. The various methods use a training set of known objects and characterize one or more properties of point clouds for the class. Then a point cloud of the unknown object is used to derive values for those one or more properties. The values of the properties for the unknown object are then matched using some measure of similarity to a class with values most similar those of the unknown object. In the following embodiments, the problem is made tractable by limiting the classes of objects to a discreet number N of classes of interest, e.g., N types of vehicles and roadside structures, or N types of tableware, or N types of animals, or N types of hand-held weapons. The unknown object then is adequately similar to one of the N classes or is rejected as not of interest.

[0067] FIG. 6A illustrates an example of a 3D point cloud 600 of an object, according to an embodiment. In some embodiments, the 3D point cloud is obtained using an optical LIDAR system, such as described above with reference to FIG. 2 through FIG. 3B. In other embodiments, the 3D point cloud is obtained from a 3D scanner other than such optical systems. FIG. 6B is illustrates a segment 607 of the 3D point cloud 600 of FIG. 6A including a point 601 and nearest neighbor points 605 around the point 601, according to an embodiment.

[0068] Surface normals 602 are depicted in FIG. 6A and approximate a normal to the surface of the object at each point of the point cloud. The surface normal 602 at each point on the point cloud 600 is approximated. In some embodiments, the surface normal 602 is approximated using a process disclosed in Klasing 2009. In some embodiments, the surface normal 602 at the point 601 is determined by initially determining the nearest neighbor points 605 around the point 601. FIG. 6B labels some of the points 605 for purposes of simplicity. As discussed in Klasing, the (x, y, z) values of the nearest neighbor points 605 are used to construct a 3x3 covariance matrix C. In one embodiment, a minimum of three points 605 are used to construct the 3x3 covariance matrix C and define a plane in 3D. In other embodiments, the minimum number of points 605 to construct a covariance matrix C should be in a range from about 10 points to about 20 points to practically deal with noise in the system and obtain a useful measurement and still produce results in real-time or near real-time, as defined below. Eigenvalues ($\lambda_1, \lambda_2, \lambda_3$) and eigenvectors (e_1, e_2, e_3) of the matrix C are then computed. These eigenvalues represent an approximation of the variance of the points 605 in each of the three directions represented by the eigenvectors. In some embodiments, the eigenvector associated with the smallest eigenvalue is used to approximate the surface normal 602 at the point 601, since this eigenvector approximates the direction of minimal variance of the points 605 and the other two eigenvector's define a plane with most of the variance.. In other embodiments, a curvature of the object surface at point 601 can be estimated using the eigenvalues as:

$$\text{Curvature} = \frac{\lambda_3}{\lambda_1 + \lambda_2 + \lambda_3} \quad (2)$$

The surface normal 602 at each point 601 can be approximated using any method beyond that disclosed in Klasing.

[0069] FIG. 6C is a block diagram 650 that illustrates an example of feature variables defined by the points 605 of a 3D point cloud within a translational and rotational invariant coordinate system, according to an embodiment. The translational and rotational invariant coordinate system is defined by the surface normal 602 at the point 601 and a plane 655 that is orthogonal to the surface normal 602 and tangent to the object surface 656 at the point 601. In some embodiments, a feature variable α is defined by a distance 662 along the plane 655 from the point 601 to the point 605. In other embodiments, a feature variable β is defined by a distance 663 perpendicular to the plane 655 from the point 601 to the point 605; or a feature variable is defined by the pair of distances α, β . In still other embodiments, a feature variable ρ is a straight distance 660 from point 601 to point 605, and is derived with Pythagorean's theorem using α and β . In still other embodiments, a feature variable θ is an angle 658 measured between the surface normal 602 and ρ . In still other embodiments, a feature variable is a surface normal 665 at the point 605 that is determined using similar methods as those used to determine the surface normal 602. In still other embodiments, ψ is an angle 664 measured between the surface normal 665 and β . In other embodiments, other feature variables are defined, such as a length of an arc with continuous smooth tangent angles between 601 and 605. A set of one or more feature variables can be used in various embodiments. Each of one or more points in the 3D point cloud has a set of values corresponding to the one or more feature variables in the feature set.

[0070] FIG. 6D is a spin image 680 that illustrates an example of a histogram of the number of points 605 of the 3D point cloud over a range of values for two of the feature variables of FIG. 6C, according to an embodiment. The horizontal axis 682 indicates the value of feature variable α in centimeters (cm). The vertical axis 684 indicates the value of feature variable β in centimeters (cm). The α - β space is divided up into a plurality of bins 686. One bin 686 is depicted in FIG. 6D for purposes of simplicity. In an example embodiment, the bins 686 have a square dimension of 1 centimeter (cm) x 1 centimeter (cm). The histogram assigns a value for each bin 686, based on the number of points 605 of the 3D point cloud with values within the α - β range defined by the bin 686. A greyscale is used to indicate the value of each bin 686, where the value of each bin 686 is represented by low values (e.g. white) through high values (e.g. black).

[0071] In some embodiments, the spin image 680 is obtained for one segment of a 3D point cloud, such as the segment 607 depicted in FIG. 6B of the point cloud 600 depicted in FIG.

6A. In an example embodiment, the segment has a dimension of approximately 0.5 meters (m). In these embodiments, multiple spin images 680 are obtained over corresponding multiple segments that span the point cloud 600. In some embodiments values of the feature variables (e.g., α , β) are obtained to define each point in each segment of the point cloud 600. In other embodiments values for the feature variables (e.g., α , β) are determined for only one or a few of the point cloud points in each segment. In these embodiments, the spin images 680 are sufficiently similar across the segments that each segment is recognizable as belonging to the same class as the object used to generate the point cloud 600. Spin images 680 are obtained for any points in the point cloud 600 based on methods disclosed in Johnson 1997.

[0072] In some embodiments, the spin image 680 for each segment is represented by a vector with a dimension equal to the number of bins 686. In these embodiments, the vector is represented by $[P_1, P_2, P_3, P_4 \dots P_B]$, where P_n represents the number of points 605 in the n^{th} bin 686 and where B is the total number of bins 686. In an example embodiment, where the α - β range of the spin image 680 is 40 centimeters (cm) by 40 centimeters (cm) and where the bin 686 has a dimension of 1 cm x 1cm, B is 1600, resulting in a 1600 dimensional vector. In some embodiments, the histogram of the spin image 680 is smoothed using a discrete Parzen window approach before the spin image 680 is represented by the vector. In other embodiments, the vector represented by the spin image 680 is normalized, to remove variability that may arise due to changes in point density.

[0073] In some embodiments, the object that generated the spin image 680 is classified in about real-time or near real-time based on a fixed number (N) of classes. For purposes of this description, real-time is based on a frame rate of the 3D scanner (e.g. LIDAR) used to capture the 3D point cloud that generates the spin image 680. An inverse of the frame rate is a time capture period during which the 3D scanner captures the 3D point cloud. In one embodiment, real-time is defined as a period within the time capture period. In one example embodiment, the frame rate is in a range from 2.5 – 10 frames per second, corresponding to a time capture period of 0.1 – 0.25 seconds (sec). This kind of time period is advantageous for identifying objects in tactical and collision avoidance applications. . Near real-time is within a factor of about ten of real-time, e.g., within about 2.5 seconds for the above example time capture periods.

[0074] Point cloud data was collected for objects in each class across a variety of perspectives and orientations. Values of the (α , β) feature variables were extracted from the point clouds for the objects in each class. Since the class membership of each point cloud is known (e.g. a box, a traffic cone, etc), the values of the feature variables were labeled with the known class. The set of labeled spin images were trained using standard linear discriminant analysis (LDA), as disclosed in Ye 2007.

[0075] In an embodiment, a dimensionality reduction step is performed on the spin image 680 vector. In one embodiment, the dimension of the spin image 680 vector is reduced from the total number of bins 686 to a reduced dimension based on a total number of classes (N). In an example embodiment, the dimension of the spin image 680 vector is reduced from the total number, 1600, of bins 686 to N-1. In some embodiments, a predetermined set of projection vectors acquired during the training of the labeled spin images in LDA are used to project a dimension of the spin image 680 vector from the total number, 1600, of bins 686 to N-1. In some embodiments, the predetermined set of projection vectors are associated with a respective set of spin images for the N classes. As disclosed in Ye 2007, in this manageable feature space, a k-d tree and NN (nearest neighbor) search is performed in order to assign class membership to any unknown object with a 3D point cloud used to generate the new spin image 680. In some embodiments, where multiple spin images 680 are obtained over the multiple segments 607 that span the point cloud 600 of the unknown object, each spin image 680 is projected to the reduced dimension (e.g. N-1) and the NN search is performed to assign class membership to each segment spin image 680. In some embodiments, the projection of the spin image 680 to the reduced dimension space and the NN search assigning class membership to the spin image 680 are performed in about real-time or near real-time. In an example embodiment, if a substantial number (e.g. 90%) of the segment spin images 680 of the point cloud 600 are assigned the same class, then the unknown object used to generate the point cloud 600 is assigned that same class. In other embodiments, learning techniques are used to identify one or more sub-sets of all spin images 680 encompassing a 3D point cloud as respective subclasses with less internal variance. In other embodiments, the dimensionality reduction step is performed on the spin image 680 vector using principal component analysis in addition to or instead of the LDA.

[0076] In some embodiments, the values of feature variables [$\alpha, \beta, \theta, \rho, \psi$] are determined for each point 605 in the point cloud 600. These feature variables are then assembled in a data matrix F:

$$F = \begin{matrix} \alpha_1 & \dots & \alpha_M \\ \beta_1 & \dots & \beta_M \\ \dots & \dots & \dots \end{matrix} \quad (3)$$

where F is an $n \times M$ matrix, n is the number of feature variables and M is the number of points 605 in the point cloud. In the illustrated embodiment, n = 5 refers to feature variables [$\alpha, \beta, \theta, \rho, \psi$]. However, n is not limited to any specific number and may be less than 5 or more than 5 or may include different feature variables than those listed above.

[0077] A mean of each feature variable [$\alpha, \beta, \theta, \rho, \psi$] over the number of M points 605 in each row of the matrix is then determined. A zero-mean data matrix \tilde{F} is assembled, by subtracting the mean of each feature variable (e.g. $\bar{\alpha}, \bar{\beta}, \bar{\theta}, \bar{\rho}, \bar{\psi}$) from a respective row of the data matrix F associated with the feature variable. The resulting zero-mean data matrix \tilde{F} is provided by:

$$\tilde{F} = \begin{matrix} \alpha_1 - \bar{\alpha} & \dots & \alpha_M - \bar{\alpha} \\ \beta_1 - \bar{\beta} & \dots & \beta_M - \bar{\beta} \\ \dots & \dots & \dots \end{matrix} \quad (4)$$

The zero-mean data matrix \tilde{F} is also an $n \times M$ matrix, as with the data matrix F. A covariance matrix C is then assembled, based on the zero-mean data matrix \tilde{F} as:

$$C = \frac{1}{N} (\tilde{F} \tilde{F}^T) = \begin{matrix} \sigma_1^2 & \dots & \sigma_1 \sigma_n \\ \dots & \dots & \dots \\ \sigma_n \sigma_1 & \dots & \sigma_n^2 \end{matrix} \quad (5)$$

where σ_n is a variance of the n^{th} feature variable from the feature variable mean over the M number of elements in the row representing the points 605 in the point cloud segment 607. The diagonal elements of the covariance matrix C represent the variance (e.g. σ_1^2) of a single feature variable, whereas non-diagonal elements of the covariance matrix C represent the covariance (e.g. $\sigma_1 \sigma_2$) of two different feature variables. The covariance matrix C is an $n \times n$ matrix, since the zero-mean data matrix \tilde{F} is an $n \times M$ matrix and matrix \tilde{F}^{-1} is an $M \times n$ matrix. In the illustrated embodiment, where each of the above five feature variables is used to assemble the covariance matrix C, the covariance matrix is a 5x5 matrix. However, the covariance matrix C can be any sized $n \times n$ matrix, where n is less than or more than 5,

depending on the number of feature variables that are used for classification purposes. In some embodiments, the covariance matrix C is constructed using the feature variables for each point, based on methods disclosed in Fehr 2012.

[0078] In some embodiments, the object that generated the point cloud used to assemble the covariance matrix C is classified about in real-time or near-real time, as defined above, based on the fixed number (N) of classes.

[0079] A large number of covariance matrices C are assembled during training, for a respective large number of point clouds associated with each of a predetermined fixed number (N) of classes. This large number of covariance matrices C are arranged in a predetermined fixed number of clusters associated with the predetermined fixed number (N) of classes, as disclosed in Salehian 2013. A center of each cluster is selected to represent the cluster and a covariance matrix C associated with the cluster center is selected to represent the class associated with the cluster. In an embodiment, this compresses the training data by reducing the number of covariance matrices to be compared with an input covariance matrix, from the initial large number of covariance matrices to the reduced number equal to the number of classes. FIG. 6E is a graph 670 that illustrates an example of a plurality of clusters 676a, 676b, 676c, 676d and corresponding cluster centers 678a, 678b, 678c, 678d, respectively, according to an embodiment. In this embodiment, the cluster centers 678a, 678b, 678c, 678d are each associated with a respective object class. For purposes of simplicity, FIG. 6E depicts four clusters associated with four classes in a two dimensional space. However, the fixed number (N) of classes may be greater than four classes, such as 10 or 100, for example, and the cluster space may have more than two dimensions in typical embodiments. The horizontal axis 672 is a first parameter and the vertical axis 674 is a second parameter. The two-dimensional space in FIG. 6E is representative of a space where the clustering occurs.

[0080] A covariance matrix C associated with each cluster center 678a, 678b, 678c, 678d is compared with test data input of a covariance matrix C. In some embodiments, a distance is computed between the input covariance matrix C and each cluster center matrix C. In an embodiment, a geodesic distance is computed between the matrices using the Jensen-Bregman LogDet divergence operation, as disclosed in Salehian 2013. FIG. 6E depicts the data input 679 associated with the input covariance matrix C. In some embodiments, a k-d tree NN search is performed to determine which of the cluster centers 678a, 678b, 678c, 678d

is the most proximate to the data input 679. In this embodiment, cluster center 678a is the most proximate to the data input 679 and thus the object used to generate the input covariance matrix C is classified based on the class associated with cluster 676a. In some embodiments, the NN search between the data input 679 and each cluster center is performed in about real-time or near-real time.

[0081] In some embodiments, a covariance matrix C is assembled for each point of the point cloud. In an embodiment, a subset of these covariance matrices C is compared with the covariance matrices associated with the cluster centers, in order to classify the object used to generate the point cloud. In another embodiment, if a majority of the subset of these covariance matrices is associated with one object class, the object is assigned this object class. In an example embodiment, if 1000 covariance matrices are assembled for 1000 points or 1000 segments of a point cloud, then a subset (e.g. 50) of the covariance matrices are compared with the covariance matrices associated with the cluster centers. In this example embodiment, if a majority (e.g. 30) of the subset of the covariance matrices is associated with one object class, then the object is classified according to this object class.

[0082] FIG. 7 is a flow chart that illustrates an example method for classifying an object defined by a 3D point cloud, according to an embodiment. Although steps are depicted in FIG. 7 as integral steps in a particular order for purposes of illustration, in other embodiments, one or more steps, or portions thereof, are performed in a different order, or overlapping in time, in series or in parallel, or are omitted, or one or more additional steps are added, or the method is changed in some combination of ways.

[0083] In step 701, a 3D point cloud representing an external surface of an object is obtained. In some embodiments, the 3D point cloud is generated using the system 200 of FIG. 2 or systems 300a, 300b of FIG. 3A and FIG. 3B. In other embodiments, the 3D point cloud is generated using any 3D scanner appreciated by one of ordinary skill in the art. In still other embodiments, the 3D point cloud is obtained from an external system and downloaded from an external source associated with the system or retrieved from storage on a local or remote storage device. In one embodiment, the 3D point cloud is obtained from an object or target in the system 200 of FIG. 2 or systems 300a, 300b of FIG. 3A and 3B that is located at a minimum range, such as 400 meters (m).

[0084] In step 702 the surface normal 602 is extracted at the point 601 on the point cloud 600. In one embodiment, the surface normal 602 is extracted using the approximation from

Klasing 2009 based on an eigenvector associated with a smallest eigenvalue that indicates a direction of minimal variance of neighbor points 605. In other embodiments, the surface normal 602 is extracted using another method. In one embodiment, the surface normal 602 is extracted for at least one point for each segment over the point cloud 600. The surface normal 602 defines the translational and rotational invariant coordinate system at the point 601 based on the surface normal 602 and the plane 655 that is orthogonal to the surface normal 602 (and, therefore, is tangent to the object surface).

[0085] In step 704, values for one or more feature variables are extracted in the translational and rotational invariant coordinate system for at least one point 605 in every segment of the point cloud 600. In an example embodiment, values for the feature variables are extracted for each point 605 in each segment 607 of the point cloud 600.

[0086] In step 705, a first classification statistic is computed for at least one point 605 in the point cloud 600 based on the extracted feature variables in step 704. In one embodiment, the first classification statistic is the spin image 680 that is based on extracted values of one or more features (α, β) from step 704. In one embodiment, the spin image 680 is obtained for each segment 607 of the point cloud 600 and thus multiple spin images 680 are obtained over the multiple segments 607 that span the point cloud 600.

[0087] In step 707, a second classification statistic is computed for each point 605 in at least one segment of the point cloud 600 based on the extracted values of the feature variables in step 704. In one embodiment, the second classification statistic is the covariance matrix C of equation 5. In some embodiments, a covariance matrix C is assembled for each point in each segment of the point cloud.

[0088] In step 709, a closest match is determined between the first classification statistic of step 705 and a set of first classification statistics for a corresponding set of N classes in about real-time or near real-time. This closest match is then used to estimate that the object generating the point cloud 600 is in a first class. In some embodiments, it is possible for the first classification statistic to not match any of the N classes. For example, the point cloud may represent an object that is not of interest, such as table setting when the objects of interest are vehicles and roadside structures. In these embodiments, in step 709, it is estimated that the object generating the point cloud is not in any of the N classes. This amounts to a negative classification (e.g. the object is not in any of the N classes) which is advantageous in various applications, particularly where the N classes encompassing objects

of interest are known. In some embodiments, it is advantageous to know if an object is “unknown” (e.g. is not in any of the pre-defined N classes) such that the associated point cloud data is stored for further analysis and training if it were deemed by the user to be of interest. In an example embodiment, a label such as “unknown weapon” is used to update the training library of point clouds so that it is available to become a known target or class upon further training.

[0089] In an embodiment, in step 709, the spin image 680 from step 705 is represented by the vector equal to the number of bins 686 in the spin image 680. Additionally, in an embodiment, in step 709, the set of labeled spin images associated with a set of known classes is trained using LDA. In one embodiment, a dimensionality reduction step is performed in about real-time or near real-time, on the spin image 680 from a high dimension (e.g. number of bins 686) to a lower dimension (e.g. N-1) using projection vectors acquired during LDA. In an embodiment, a k-d tree and NN search is performed in about real-time or near-real time, to assign the object to membership in a first class. In an embodiment, the number (N) of the set of known classes and consequently the number of the set of first classification statistics is limited to a maximum threshold. In one example embodiment, N is less than or equal to 10. In another example embodiment, N is less than or equal to 100. In one embodiment, step 709 is performed in about real-time or near real-time as a result of the dimensional reduction of the spin image 680 vector and resulting k-d tree NN search in the reduced dimensional space.

[0090] In step 711, a closest match is determined between the second classification statistic of step 707 and a set of second classification statistics for a set of N classes in about real-time or in near real-time. This closest match is then used to estimate that the object generating the point cloud 600 is in a second class. In some embodiments, it is possible for the second classification statistic to not match any of the N classes. For example, the point cloud may represent an object that is not of interest, such as table setting when the objects of interest are vehicles and roadside structures. In these embodiments, in step 711, it is estimated that the object generating the point cloud is not in any of the N classes. This amounts to a negative classification (e.g. the object is not in any of the N classes) which is advantageous in various applications, particularly where the N classes encompassing objects of interest are known.

[0091] In an embodiment, in step 711, prior to a data capture period, a large number of covariance matrices C are assembled and associated with a fixed number (N) of classes. This

large number of covariance matrices C are arranged in a fixed number of clusters based on the fixed number (N) of classes, as disclosed in Salehian 2013. A center of each cluster is selected to represent each cluster and a covariance matrix associated with the cluster center is selected to represent each class. In an embodiment, in step 711, a distance is computed in about real-time or in near-real-time between a covariance matrix C from test data input and the covariance matrices associated with each class. In an example embodiment, k-d tree NN searching is used to determine the nearest neighbor to the covariance matrix C in about real-time or near real-time. The object is assigned to a second class associated with the closest cluster center. In an embodiment, the number (N) of the set of known classes and consequently the number of cluster centers is limited to a maximum threshold. In one example embodiment, N is less than or equal to 10. In another example embodiment, N is less than or equal to 100. In one embodiment, step 711 is performed in about real-time or near real-time as a result of the reduction of the covariance matrices C to the fixed number of clusters and resulting k-d tree NN search based on this fixed number of clusters.

[0092] In step 713, the first class from step 709 is compared to the second class from step 711. If the first class is the same as the second class, then the method moves to step 715 where the object is assigned to the first class. If the first class is not the same as the second class, then the method moves to step 717. In other embodiments, if step 709 estimates that the object is not in any of the N classes and if step 711 estimates that the object is not in any of the N classes, then the method does not assign any of the N classes to the object. In one embodiment, the method outputs on a display that the object is not in any of the N classes. In another embodiment, the method operates a device based on the object not being in any of the N classes.

[0093] In some embodiments, where one of step 709 or step 711 concludes that the object is not in any of the N classes and the other of step 709 or step 711 concludes that the object is in one of the N classes, step 713 concludes that the classifications in steps 709 and 711 do not agree and thus the method moves to step 717. However, in these embodiments, as discussed below, in step 717, a closest fit is performed between the point cloud and a model point cloud for only the single class identified in one of step 709 or step 711 that identified the object in one of the N classes.

[0094] In step 717, a third classification statistic is computed in about real-time or near real-time for each point cloud 600 based on a closest fit between the point cloud 600 and one or

more predetermined occluded or non-occluded model point clouds for the first and second classes only. None of the other N-2 classes is considered. In one embodiment, step 717 is performed in about real-time or near real-time as a result of the restriction of the closest fit to only model point clouds for the first class and second class.

[0095] This offers an advantage of saving valuable time for this method which can be the most time consuming classifier, e.g., involving the most computations. The number of computations is related to the number of points in the point cloud times the number of models which is N times the number of surfaces to define each of the N classes. In one embodiment, the third classification statistic is an iterative closest point (ICP) classification statistic as disclosed in Besl 1992. Step 717 is used to determine in about real-time or near real-time which class (e.g. first class or second class) is a better fit for the object generating the point cloud 600. In some embodiments, in step 717 the closest fit is only performed between the point cloud 600 and predetermined model point clouds for the first and second classes.

[0096] In some embodiments, in step 717, model point clouds are obtained for the first and second class. In one embodiment, the model point cloud is obtained occluding those portions of the point cloud that are not observable from each perspective. FIG. 10A is a block diagram that illustrates an example of an occluded model 1000 of the object based on the 3D point cloud 400 of the object in FIG. 4, according to an embodiment. In this example embodiment, occluded model 1002a is obtained on an elevation view of the handle portion of the cup. Those portions of the cup (e.g. far side of the cup) not observable in that perspective are omitted in the occluded model 1002a. Similarly, occluded models 1002b, 1002c are obtained on a plan view of the rim portion of the cup near the top where the diameter is greater and the handle is present and near the bottom of the cup, where the diameter is less and there is no handle.. Those portions of the cup (e.g. back side of the cup) are not observable in these perspectives and are omitted in the occluded models 1002b, 1002c. In a similar manner, occluded model point clouds are obtained of the first and second classes, using a 3D point cloud of the first and second classes. In some embodiments, the occluded model point clouds are generated using a master point cloud of the object which is randomly rotated and then converted to spherical coordinates. In this coordinate system, the closest point in range was selected for all occupied angle patches. Each angle patch has an empirically chosen extent (e.g. 0.001 degrees x. 0.001 degrees) which is converted back to Cartesian coordinates to constitute an instance of possible occlusion. These sub-models of

model 1000 would be used in step 717 if one class of the two different classes determined in steps 709 and 711 were a cup and the second class were different, e.g., a platter.

[0097] In some embodiments, in step 717, the closest fit is performed between a test input point cloud and the model point clouds associated with the first and second classes, in order to determine whether or not the object should be classified in the first or second class. FIG. 10B is a block diagram that illustrates an example of a 2D test point cloud 1010 of an object, according to an embodiment. The 2D point cloud 1010 is a test input point cloud that will be compared to each model point cloud of FIG. 10A, using a closest fit. The 2D point cloud is rotated and translated to best fit each model for the class. FIGS. 10C-10E are block diagrams that illustrate an example of a closest fit between the 2D point cloud 1010 of FIG. 10B and each occluded model 1002a, 1002b, 1002c of FIG. 10A, according to an embodiment. In each closest fit, the test input point cloud (e.g. 1010) is rotated, translated and/or scaled relative such that a minimum distance is obtained between points on the model point cloud and points on the test input point cloud, as disclosed in Besl 1992. The fit with the smallest minimum distance is designated as the closest fit. If the closest fit is too large, e.g., the mean square distance between points in the test input point cloud and points in the model point cloud for a minimum ratio of closest points between the point clouds is above a threshold square distance, then the object is considered not to belong to the class. In an example embodiment, if the mean square distance between points in the test input point cloud and points in the model point cloud for the top 90% of closest points between the point clouds is above 2cm², the object is considered not to belong to the class.

[0098] In the example embodiment of FIG. 10C, the best fit between the model point cloud 1002a and model point 1010 results in a relatively large average distance (or root mean square distance) between the point clouds. In the example embodiment of FIG. 10D, the best fit between the model point cloud 1002b and model point 1010 results in a smaller average distance or root mean square distance (relative to FIG. 10C) between the point clouds. In the example embodiment of FIG. 10E, the best fit between the model point cloud 1002c and model point 1010 results in the smallest average or root mean square distance between the point clouds. If this distance is acceptably small, the 2D point cloud 1010 is successfully classified as a cup.. Similarly the point cloud 1010 would be compared to a model of the second class of objects determined in step 711, such as a platter with several occluded sub-

models. If the best fit of the platter sub-models has a larger average or root mean square distance, then the object would be classified as a cup and not a platter.

[0099] In some embodiments, where one of step 709 or step 711 concludes that the object is not in any of the N classes and the other of step 709 or step 711 concludes that the object is in one of the N classes, then the closest fit in step 717 is performed between the test input point cloud and the model point cloud for only the single class associated with step 709 or step 711 that concludes the object is in one of the N classes. In these embodiments, if the smallest average or root mean square distance is acceptably small, then the object is classified with the same classification as the single class associated with step 709 or step 711 that concludes the object is in one of the N classes. In these embodiments, if the smallest average or root mean square distance is not acceptably small, then the object is not classified in any of the N classes.

[0100] In step 719, the object is assigned to the first or second class, depending on which class results in the closest fit with the input point cloud in step 717.

[0101] In step 721, after the object is assigned a class, a device is operated based on the assigned object class. In some embodiments, this involves presenting an image on a display device that indicates the assigned class or information based on the assigned class. In some embodiments, this involves communicating, to the device, data that identifies the assigned object class. In other embodiments, this involves impelling a projectile along a trajectory to the object. In an example embodiment, the projectile is a missile. In some embodiments, the device is a controlled vehicle, and based on the object, the vehicle is controlled to avoid collision with the object or controlled to not avoid collision with the object.

6. Example Embodiments

[0102] In these example embodiments, the LIDAR system used components illustrated above to produce simultaneous up and down chirp transmitted signals. This system is commercially available as HRS-3D from BLACKMORE SENSORS AND ANALYTICS, INC.TM of Bozeman.

[0103] FIG. 8A is a photograph 800 that illustrates an example of a top perspective view of an experimental setup of the system 200 of FIG. 2 or system 300 of FIG. 3A or FIG. 3B, according to embodiments. In an embodiment, the objects 802 are spaced apart from the LIDAR system in a test lab 804 by a range 806. In an example embodiment, the range 806 is between 400 meters (m) and 2000 meters (m), for example. FIG. 8B is a block diagram that

that illustrates an example of objects 802 used in the experimental setup depicted in FIG. 8A, according to an embodiment. These objects 802 include a cone 802a, a box 802b, a pipe 802c, a pistol 802d and a box 802e, representing $N = 5$ object classes of interest. A point cloud dissimilar from all these objects, using any similarity measure and an associated threshold similarity value, is not an object of interest. Table 1 below depicts a matrix where the actual object class (vertical column) is compared with the predicted object class (horizontal row) using the method 700 of FIG. 7. The sum of all values in a row is 100%. In an example embodiment, the method 700 predicted that a 3D point cloud of the cone is in the cone class 100% of the time; that a 3D point cloud of the cube is in the cube class 100% of the time; that a 3D point cloud of the pipe is in the pipe class 96% of the time and in the pistol class 4% of the time; that a 3D point cloud of the box is in the cube class 2% of the time, in the pipe class 2% of the time and is in the box class 96% of the time; and predicted that a 3D point cloud of the pistol is in the cone class 1% of the time, in the pipe class 10% of the time and is in the pistol class 89% of the time.

Table 1

	Cone	Cube	Pipe	Box	Pistol
Cone	100				
Cube		100			
Pipe			96		4
Box		2.1	2.1	95.8	
Pistol	1.1		9.5		89.4

[0104] FIG. 9A is a set of point clouds 900 for a set of objects, according to an embodiment. FIG. 9B is a block diagram that illustrates an example of the set of objects 902 used to obtain the set of point clouds 900 of FIG. 9A according to an embodiment. The objects 902 are more similar in shape and size than the objects 802. For example, the three toy guns 902a, 902c, 902g, look similar. The objects 902 include a toy gun 902a (“Sledge” in Table 2), a box 902b, a toy gun 902c (“Mod” in Table 2), a cone 902d, a pipe 902e, a cube 902f, and a toy gun 902g (“Roto” in Table 2).

[0105] In one embodiment, the objects 902 are positioned in an experimental setup that is similar to that depicted in FIG. 8A. Table 2 below depicts a matrix where the actual object class (vertical column) is compared with the predicted object class (horizontal row), according to the method 700 of FIG. 7. Again the sum of all columns in each row is 100%. In an example embodiment, the method 700 predicted that a 3D point cloud of the cone is in

the cone class 100% of the time; that a 3D point cloud of the cube is in the cube class 100% of the time; that a 3D point cloud of the pipe is in the pipe class 100% of the time; that a 3D point cloud of the mod toy gun is in the cube class 1% of the time, in the mod class 92% of the time, in the roto class 6% of the time and is in the sledge class 1% of the time; that a 3D point cloud of the roto toy gun is in the pipe class 3% of the time and in the roto class 97% of the time; that a 3D point cloud of the box is in the pipe class 5% of the time, in the mod class 1% of the time, in the box class 93% of the time and is in the sledge class 1% of the time; and that a 3D point cloud of the sledge toy gun is in the pipe class 3% of the time, in the box class 10% of the time and is in the sledge class 87% of the time.

Table 2

	Cone	Cube	Pipe	Mod	Roto	Box	Sledge
Cone	100						
Cube		100					
Pipe			100				
Mod		1		92	6		1
Roto			3		97		
Box			5	1		93	1
Sledge			3			10	87

Based on the results of Table 2, the method 700 maintained strong class predictions, despite a more challenging object set that are more similar in size and shape.

7. Computational Hardware Overview

[0106] FIG. 11 is a block diagram that illustrates a computer system 1100 upon which an embodiment of the invention may be implemented. Computer system 1100 includes a communication mechanism such as a bus 1110 for passing information between other internal and external components of the computer system 1100. Information is represented as physical signals of a measurable phenomenon, typically electric voltages, but including, in other embodiments, such phenomena as magnetic, electromagnetic, pressure, chemical, molecular atomic and quantum interactions. For example, north and south magnetic fields, or a zero and non-zero electric voltage, represent two states (0, 1) of a binary digit (bit). Other phenomena can represent digits of a higher base. A superposition of multiple simultaneous quantum states before measurement represents a quantum bit (qubit). A sequence of one or more digits constitutes digital data that is used to represent a number or

code for a character. In some embodiments, information called analog data is represented by a near continuum of measurable values within a particular range. Computer system 1100, or a portion thereof, constitutes a means for performing one or more steps of one or more methods described herein.

[0107] A sequence of binary digits constitutes digital data that is used to represent a number or code for a character. A bus 1110 includes many parallel conductors of information so that information is transferred quickly among devices coupled to the bus 1110. One or more processors 1102 for processing information are coupled with the bus 1110. A processor 1102 performs a set of operations on information. The set of operations include bringing information in from the bus 1110 and placing information on the bus 1110. The set of operations also typically include comparing two or more units of information, shifting positions of units of information, and combining two or more units of information, such as by addition or multiplication. A sequence of operations to be executed by the processor 1102 constitutes computer instructions.

[0108] Computer system 1100 also includes a memory 1104 coupled to bus 1110. The memory 1104, such as a random access memory (RAM) or other dynamic storage device, stores information including computer instructions. Dynamic memory allows information stored therein to be changed by the computer system 1100. RAM allows a unit of information stored at a location called a memory address to be stored and retrieved independently of information at neighboring addresses. The memory 1104 is also used by the processor 1102 to store temporary values during execution of computer instructions. The computer system 1100 also includes a read only memory (ROM) 1106 or other static storage device coupled to the bus 1110 for storing static information, including instructions, that is not changed by the computer system 1100. Also coupled to bus 1110 is a non-volatile (persistent) storage device 1108, such as a magnetic disk or optical disk, for storing information, including instructions, that persists even when the computer system 1100 is turned off or otherwise loses power.

[0109] Information, including instructions, is provided to the bus 1110 for use by the processor from an external input device 1112, such as a keyboard containing alphanumeric keys operated by a human user, or a sensor. A sensor detects conditions in its vicinity and transforms those detections into signals compatible with the signals used to represent information in computer system 1100. Other external devices coupled to bus 1110, used

primarily for interacting with humans, include a display device 1114, such as a cathode ray tube (CRT) or a liquid crystal display (LCD), for presenting images, and a pointing device 1116, such as a mouse or a trackball or cursor direction keys, for controlling a position of a small cursor image presented on the display 1114 and issuing commands associated with graphical elements presented on the display 1114.

[0110] In the illustrated embodiment, special purpose hardware, such as an application specific integrated circuit (IC) 1120, is coupled to bus 1110. The special purpose hardware is configured to perform operations not performed by processor 1102 quickly enough for special purposes. Examples of application specific ICs include graphics accelerator cards for generating images for display 1114, cryptographic boards for encrypting and decrypting messages sent over a network, speech recognition, and interfaces to special external devices, such as robotic arms and medical scanning equipment that repeatedly perform some complex sequence of operations that are more efficiently implemented in hardware.

[0111] Computer system 1100 also includes one or more instances of a communications interface 1170 coupled to bus 1110. Communication interface 1170 provides a two-way communication coupling to a variety of external devices that operate with their own processors, such as printers, scanners and external disks. In general the coupling is with a network link 1178 that is connected to a local network 1180 to which a variety of external devices with their own processors are connected. For example, communication interface 1170 may be a parallel port or a serial port or a universal serial bus (USB) port on a personal computer. In some embodiments, communications interface 1170 is an integrated services digital network (ISDN) card or a digital subscriber line (DSL) card or a telephone modem that provides an information communication connection to a corresponding type of telephone line. In some embodiments, a communication interface 1170 is a cable modem that converts signals on bus 1110 into signals for a communication connection over a coaxial cable or into optical signals for a communication connection over a fiber optic cable. As another example, communications interface 1170 may be a local area network (LAN) card to provide a data communication connection to a compatible LAN, such as Ethernet. Wireless links may also be implemented. Carrier waves, such as acoustic waves and electromagnetic waves, including radio, optical and infrared waves travel through space without wires or cables. Signals include man-made variations in amplitude, frequency, phase, polarization or other physical properties of carrier waves. For wireless links, the communications interface 1170

sends and receives electrical, acoustic or electromagnetic signals, including infrared and optical signals, that carry information streams, such as digital data.

[0112] The term computer-readable medium is used herein to refer to any medium that participates in providing information to processor 1102, including instructions for execution. Such a medium may take many forms, including, but not limited to, non-volatile media, volatile media and transmission media. Non-volatile media include, for example, optical or magnetic disks, such as storage device 1108. Volatile media include, for example, dynamic memory 1104. Transmission media include, for example, coaxial cables, copper wire, fiber optic cables, and waves that travel through space without wires or cables, such as acoustic waves and electromagnetic waves, including radio, optical and infrared waves. The term computer-readable storage medium is used herein to refer to any medium that participates in providing information to processor 1102, except for transmission media.

[0113] Common forms of computer-readable media include, for example, a floppy disk, a flexible disk, a hard disk, a magnetic tape, or any other magnetic medium, a compact disk ROM (CD-ROM), a digital video disk (DVD) or any other optical medium, punch cards, paper tape, or any other physical medium with patterns of holes, a RAM, a programmable ROM (PROM), an erasable PROM (EPROM), a FLASH-EPROM, or any other memory chip or cartridge, a carrier wave, or any other medium from which a computer can read. The term non-transitory computer-readable storage medium is used herein to refer to any medium that participates in providing information to processor 1102, except for carrier waves and other signals.

[0114] Logic encoded in one or more tangible media includes one or both of processor instructions on a computer-readable storage media and special purpose hardware, such as ASIC 1120.

[0115] Network link 1178 typically provides information communication through one or more networks to other devices that use or process the information. For example, network link 1178 may provide a connection through local network 1180 to a host computer 1182 or to equipment 1184 operated by an Internet Service Provider (ISP). ISP equipment 1184 in turn provides data communication services through the public, world-wide packet-switching communication network of networks now commonly referred to as the Internet 1190. A computer called a server 1192 connected to the Internet provides a service in response to

information received over the Internet. For example, server 1192 provides information representing video data for presentation at display 1114.

[0116] The invention is related to the use of computer system 1100 for implementing the techniques described herein. According to one embodiment of the invention, those techniques are performed by computer system 1100 in response to processor 1102 executing one or more sequences of one or more instructions contained in memory 1104. Such instructions, also called software and program code, may be read into memory 1104 from another computer-readable medium such as storage device 1108. Execution of the sequences of instructions contained in memory 1104 causes processor 1102 to perform the method steps described herein. In alternative embodiments, hardware, such as application specific integrated circuit 1120, may be used in place of or in combination with software to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware and software.

[0117] The signals transmitted over network link 1178 and other networks through communications interface 1170, carry information to and from computer system 1100. Computer system 1100 can send and receive information, including program code, through the networks 1180, 1190 among others, through network link 1178 and communications interface 1170. In an example using the Internet 1190, a server 1192 transmits program code for a particular application, requested by a message sent from computer 1100, through Internet 1190, ISP equipment 1184, local network 1180 and communications interface 1170. The received code may be executed by processor 1102 as it is received, or may be stored in storage device 1108 or other non-volatile storage for later execution, or both. In this manner, computer system 1100 may obtain application program code in the form of a signal on a carrier wave.

[0118] Various forms of computer readable media may be involved in carrying one or more sequence of instructions or data or both to processor 1102 for execution. For example, instructions and data may initially be carried on a magnetic disk of a remote computer such as host 1182. The remote computer loads the instructions and data into its dynamic memory and sends the instructions and data over a telephone line using a modem. A modem local to the computer system 1100 receives the instructions and data on a telephone line and uses an infra-red transmitter to convert the instructions and data to a signal on an infra-red a carrier wave serving as the network link 1178. An infrared detector serving as communications

interface 1170 receives the instructions and data carried in the infrared signal and places information representing the instructions and data onto bus 1110. Bus 1110 carries the information to memory 1104 from which processor 1102 retrieves and executes the instructions using some of the data sent with the instructions. The instructions and data received in memory 1104 may optionally be stored on storage device 1108, either before or after execution by the processor 1102.

[0119] FIG. 12 illustrates a chip set 1200 upon which an embodiment of the invention may be implemented. Chip set 1200 is programmed to perform one or more steps of a method described herein and includes, for instance, the processor and memory components described with respect to FIG. 11 incorporated in one or more physical packages (e.g., chips). By way of example, a physical package includes an arrangement of one or more materials, components, and/or wires on a structural assembly (e.g., a baseboard) to provide one or more characteristics such as physical strength, conservation of size, and/or limitation of electrical interaction. It is contemplated that in certain embodiments the chip set can be implemented in a single chip. Chip set 1200, or a portion thereof, constitutes a means for performing one or more steps of a method described herein.

[0120] In one embodiment, the chip set 1200 includes a communication mechanism such as a bus 1201 for passing information among the components of the chip set 1200. A processor 1203 has connectivity to the bus 1201 to execute instructions and process information stored in, for example, a memory 1205. The processor 1203 may include one or more processing cores with each core configured to perform independently. A multi-core processor enables multiprocessing within a single physical package. Examples of a multi-core processor include two, four, eight, or greater numbers of processing cores. Alternatively or in addition, the processor 1203 may include one or more microprocessors configured in tandem via the bus 1201 to enable independent execution of instructions, pipelining, and multithreading. The processor 1203 may also be accompanied with one or more specialized components to perform certain processing functions and tasks such as one or more digital signal processors (DSP) 1207, or one or more application-specific integrated circuits (ASIC) 1209. A DSP 1207 typically is configured to process real-world signals (e.g., sound) in real time independently of the processor 1203. Similarly, an ASIC 1209 can be configured to perform specialized functions not easily performed by a general purposed processor. Other specialized components to aid in performing the inventive functions described herein include

one or more field programmable gate arrays (FPGA) (not shown), one or more controllers (not shown), or one or more other special-purpose computer chips.

[0121] The processor 1203 and accompanying components have connectivity to the memory 1205 via the bus 1201. The memory 1205 includes both dynamic memory (e.g., RAM, magnetic disk, writable optical disk, etc.) and static memory (e.g., ROM, CD-ROM, etc.) for storing executable instructions that when executed perform one or more steps of a method described herein. The memory 1205 also stores the data associated with or generated by the execution of one or more steps of the methods described herein.

8. Alterations, Extensions and Modifications

[0122] In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. Throughout this specification and the claims, unless the context requires otherwise, the word “comprise” and its variations, such as “comprises” and “comprising,” will be understood to imply the inclusion of a stated item, element or step or group of items, elements or steps but not the exclusion of any other item, element or step or group of items, elements or steps. Furthermore, the indefinite article “a” or “an” is meant to indicate one or more of the item, element or step modified by the article. As used herein, unless otherwise clear from the context, a value is “about” another value if it is within a factor of two (twice or half) of the other value. While example ranges are given, unless otherwise clear from the context, any contained ranges are also intended in various embodiments. Thus, a range from 0 to 10 includes the range 1 to 4 in some embodiments.

9. References

Adany, P., C. Allen, and R. Hui, "Chirped Lidar Using Simplified Homodyne Detection," *Jour. Lightwave Tech.*, v. 27 (16), 15 August, 2009.

P.J. Besl and N.D. McKay, "Method for registration of 3-D shapes," 1992, vol. 1611, pp. 586-606.

D. Fehr, A. Cherian, R. Sivalingam, S. Nickolay, V. Moreallas, and N. Papanikolopoulos, "Compact Covariance descriptors in 3D point clouds for object recognition," presented at the Robotics and Automation (ICRA), 2012 IEEE International Conference, pp. 1793-1798.

Hui, R., C. Allen, and P. Adany, "Coherent detection scheme for FM Chirped laser RADAR," US patent 7,742,152, 22 June 2010.

Andrew Johnson, "Spin-Images: A Representation for 3-D Surface Matching," doctoral dissertation, tech. report CMU-RI-TR-97-47, Robotics Institute, Carnegie Mellon University, August, 1997.

Kachelmyer, A.L., "Range-Doppler Imaging with a Laser Radar," *The Lincoln Laboratory Journal*, v. 3. (1), 1990.

K. Klasing, D. Althoff, D. Wollher and M. Buss, "Comparison of Surface Normal Estimation Methods for Range Sensing Applications," in Proceedings of the 2009 IEEE International Conference on Robotics and Automation, Piscataway, NJ, USA, 2009, p. 1977-1982.

H. Salehian, G. Cheng, B.C. Vemuri and J. Ho, "Recursive Estimation of the Stein Center of SPD Matrices and Its Applications," in 2013 IEEE International Conference on Computer Vision (ICCV), 2013, pp. 1793-1800.

J. Ye, "Least Squares Linear Discriminant Analysis," Proceedings of the 24th International Conference on Machine Learning, p. 1087-1093.

CLAIMS

What is claimed is:

1. A method implemented on a processor comprising:
 - obtaining a 3D point cloud representing an external surface of an object;
 - extracting a surface normal at a point on the 3D point cloud to define a translational and rotational invariant coordinate system at the point;
 - extracting values of at least one feature variable in the coordinate system to define at least one point in the 3D point cloud;
 - computing a first classification statistic for the at least one point in the 3D point cloud based on values of the at least one feature variable;
 - computing a second classification statistic for the at least one point in the 3D point cloud based on values of the at least one feature variable;
 - determining a closest match between the first classification statistic and a first one of a set of N first classification statistics corresponding to a set of N classes of a first classifier to estimate that the object belongs in a first class;
 - determining a closest match between the second classification statistic and a second one of a set of second classification statistics corresponding to the N set of classes of a second classifier to estimate that the object belongs in a second class;
 - computing a third classification statistic for the at least one point in the 3D point cloud based on a closest fit between the 3D point cloud and model point clouds for only the first class and the second class of a third classifier, if the first class does not correspond to the second class;
 - assigning the object to one of the first class and second class based on the closest fit of the third classifier within near real time from receiving the 3D point cloud; and
 - operating a device based on the one class that the object is assigned to.
2. A method as recited in claim 1, wherein the first classification statistic is based on a spin image comprising a histogram of a number of points in the 3D point cloud in each of a number of bins over a range of values of the at least one feature variable.

3. A method as recited in claim 2, wherein the determining the closest match comprises:

converting the spin image into a vector comprising the number of points for each bin, wherein a dimension of the vector is the number of bins;

compressing the dimension of the vector from the number of bins to a reduced dimension based on the set of classes as the first classification statistic;

performing a search to determine a nearest neighbor between the reduced dimension vector and a predetermined set of reduced dimension vectors associated with a respective set of spin images for the N classes retrieved from a computer-readable medium; and

estimating the object belongs in the first class based on the reduced dimension vector associated with the first class being the nearest neighbor.

4. A method as recited in claim 1, wherein the second classification statistic is based on a covariant descriptor comprising a matrix with diagonal elements based on a variance of one feature variable from a mean of the feature variable over the 3D point cloud and non-diagonal elements based on a variance of different feature variables from a respective mean of the different feature variables over the 3D point cloud.

5. A method as recited in claim 4, wherein the determining the closest match between the covariant descriptor and the set of covariant descriptors comprises:

retrieving from a computer-readable medium a predetermined set of clusters corresponding to the respective set of classes including a cluster center for each respective cluster and a covariant descriptor including a matrix for each cluster center; and

performing a search of a nearest neighbor between the matrix of the 3D point cloud and the matrix for each cluster center based on computing a distance between the matrix of the 3D point cloud and the matrix for each cluster center; and

estimating the object belongs in the second class based on the matrix for the cluster center associated with the second class being the nearest neighbor.

6. A method as recited in claim 1, wherein the third classification statistic is based on an iterative closest point (ICP) and wherein the computing the ICP comprises:
 - at least one of rotating, translating and scaling the 3D point cloud and model point cloud for the first class until a first minimum distance is obtained between the points of the 3D point cloud and points on a predetermined model point cloud for the first class;
 - at least one of rotating, translating and scaling the 3D point cloud and model point cloud for the second class until a second minimum distance is obtained between the points of the 3D point cloud and points of a predetermined model point cloud for the second class; and

determining the closest fit based on a smaller of the first minimum distance and the second minimum distance.
7. A method as recited in claim 6, wherein the predetermined model point clouds of the first and second class are determined by:
 - obtaining an initial point cloud for the first class and an initial point cloud for the second class;
 - rotating the initial point cloud for the first class and the initial point cloud for the second class to a plurality of perspectives;
 - determining a subset of the initial point cloud for the first class and the initial point cloud for the second class at each perspective that would be observed due to occlusion effects.
8. A method as recited in claim 1, wherein extracting the surface normal comprises:
 - determining a plurality of nearest neighbors to the point on the 3D point cloud;
 - determining a direction of minimal variance of the plurality of nearest neighbors;and
approximating the surface normal based on the direction of minimal variance.
9. A method as recited in claim 1, wherein N is less than about 100.

10. A method as recited in claim 1, wherein operating the device comprises at least one of outputting an image on a display and impelling a projectile along a trajectory to the object.
11. A method as recited in claim 1, further comprising assigning the object to the first class if the first class corresponds to the second class.
12. A method as recited in claim 1, wherein the assigning step is performed within real-time upon obtaining the 3D point-cloud.

13. A computer-readable medium carrying one or more sequences of instructions, wherein execution of the one or more sequences of instructions by one or more processors causes the one or more processors to perform the steps of:
 - obtaining a 3D point cloud representing an external surface of an object;
 - extracting a surface normal at a point on the 3D point cloud to define a translational and rotational invariant coordinate system at the point;
 - extracting values of at least one feature variable in the coordinate system to define at least one point in the 3D point cloud;
 - computing a first classification statistic for the at least one point in the 3D point cloud based on values of the at least one feature variable;
 - computing a second classification statistic for the at least one point in the 3D point cloud based on values of the at least one feature variable;
 - determining a closest match between the first classification statistic and a first one of a set of N first classification statistics corresponding to a set of N classes of a first classifier to estimate that the object belongs in a first class;
 - determining a closest match between the second classification statistic and a second one of a set of second classification statistics corresponding to the N set of classes of a second classifier to estimate that the object belong in a second class;
 - computing a third classification statistic for each point in the 3D point cloud based on a closest fit between the 3D point cloud and model point clouds for only the first class and the second class of a third classifier, if the first class does not correspond to the second class;
 - assigning the object to one of the first class and second class based on the closest fit of the third classifier within near real time from receiving the 3D point cloud; and
 - operating a device based on one class that the object is assigned to.
14. A computer-readable medium of claim 13, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to further perform the step of assigning the object to the first class if the first class corresponds to the second class.

15. An apparatus comprising:
 - at least one processor; and
 - at least one memory including one or more sequences of instructions, the at least one memory and the one or more sequences of instructions configured to, with the at least one processor, cause the apparatus to perform at least the following,
 - obtaining a 3D point cloud representing an external surface of an object;
 - extracting a surface normal at a point on the 3D point cloud to define a translational and rotational invariant coordinate system at the point;
 - extracting values of at least one feature variable in the coordinate system to define at least one point in the 3D point cloud;
 - computing a first classification statistic for the at least one point in the 3D point cloud based on values of the at least one feature variable;
 - computing a second classification statistic for the at least one point in the 3D point cloud based on values of the at least one feature variable;
 - determining a closest match between the first classification statistic and a first one of a set of N first classification statistics corresponding to a set of N classes of a first classifier to estimate that the object belongs in a first class;
 - determining a closest match between the second classification statistic and a second one of a set of second classification statistics corresponding to the N set of classes of a second classifier to estimate that the object belongs in a second class;
 - computing a third classification statistic for each point in the 3D point cloud based on a closest fit between the 3D point cloud and model point clouds for only the first class and second class of a third classifier, if the first class does not correspond to the second class;
 - assigning the object to one of the first class and second class based on the closest fit of the third classifier within near real time from receiving the 3D point cloud and
 - causing a device to be operated based on the one class that the object is assigned to.

16. The apparatus of claim 15, wherein the at least one memory and the one or more sequences of instructions are configured to, with the at least one processor, cause the apparatus to assign the object to the first class if the first class corresponds to the second class.

17. A system comprising:

the apparatus as claimed in claim 15;
a sensor configured to direct a transmitted signal outside the apparatus

and further configured to receive a return signal backscattered from the object

illuminated by the transmitted signal;

wherein the at least one memory and the one or more sequences of instructions are configured to, with the at least one processor, cause the processor to receive an electrical signal from the sensor, wherein the 3D point cloud is obtained based on the electrical signal.

18. A system as recited in claim 17, wherein the sensor comprises a laser source to direct the transmitted signal and an optical detector to receive the return signal.

19. An apparatus comprising:
 - a laser source configured to provide an optical signal that comprises a pulse in an optical frequency band;
 - a splitter configured to receive the signal and produce an optical transmitted signal and an optical reference signal;
 - an optical coupler configured to direct the transmitted signal outside the apparatus and to receive any return signal backscattered from an object illuminated by the transmitted signal;
 - an optical detector disposed to receive the reference signal and the return signal;
 - a processor configured to perform the steps of:
 - receiving an electrical signal from the optical detector;
 - obtaining a 3D point cloud representing an external surface of the object;
 - extracting a surface normal at a point on the 3D point cloud to define a translational and rotational invariant coordinate system at the point;
 - extracting feature variables in the coordinate system to define each point in the 3D point cloud;
 - computing a first classification statistic for each point in the 3D point cloud based on at least one feature variable;
 - computing a second classification statistic for each point in the 3D point cloud based on at least one feature variable;
 - determining a closest match between the first classification statistic and a first one of a set of N first classification statistics corresponding to a set of N classes of a first classifier to estimate that the object belongs in a first class;
 - determining a closest match between the second classification statistic and a second one of a set of second classification statistics corresponding to the set of N classes of a second classifier to estimate that the object belongs in a second class;
 - computing a third classification statistic for each point in the 3D point cloud based on a closest fit between the 3D point cloud and model point clouds for only the first class and the second class of a third classifier, if the first class does not correspond to the second class;

assigning the object to one of the first class and second class based on the closest fit of the third classifier within near real time from receiving the 3D point cloud; and
causing a device to be operated based on one class that the object is assigned to.

20. An apparatus as in claim 19, where the processor is further configured to assign the object to the first class if the first class corresponds to the second class.

FIG. 1A




FIG. 1B




FIG. 2




FIG. 3A




FIG. 3B






FIG. 4



K-D TREE 550

FIG. 5B




FIG. 6A




FIG. 6B




FIG. 6C




FIG. 6D




FIG. 6E




FIG. 7






FIG. 8A**FIG. 8B**



FIG. 9A**FIG. 9B**



FIG. 10A**FIG. 10B**



FIG. 10C**FIG. 10D**


FIG. 10E

FIG. 11





FIG. 12

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2017/062721

A. CLASSIFICATION OF SUBJECT MATTER
 IPC(8) - G01C 3/00; G06K 9/62; H04N 7/18 (2018.01)
 CPC - G01S 17/89; G06K 9/00791; G06T 17/075 (2018.01)

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

See Search History document

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

USPC - 348/142; 356/5.09; 382/106 (keyword delimited)

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

See Search History document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	LU et al. "Recognizing Objects in 3D Point Clouds with Multi-Scale Local Features," Sensors 2014, 14, 24156-24173; doi:10.3390/s141224156, 15 December 2014 [retrieved on 2018-01-11]. Retrieved from the Internet: <URL: www.mdpi.com/1424-8220/14/12/24156/pdf>, pp 24156-24173.	1-3, 6-9, 11-16 — 4, 5, 10, 17-20
Y	Cao et al. "Lidar Signal Depolarization by Solid Targets and its Application to Terrain Mapping and 3D Imagery," Defence R&D, Contract Report DRDC Valcartier CR 2011-236 March 2011 [retrieved on 2018-01-11]. Retrieved from the Internet: <URL: http://cradpdf.drdc-rddc.gc.ca/PDFS/unc126/p537764_A1b.pdf> pp 1-74	4, 5, 17-20
Y	US 2012/0274922 A1 (HODGE) 01 November 2012 (01.11.2012) entire document	10
A	US 2010/0188504 A1 (DIMSDALE et al) 29 July 2010 (29.07.2010) entire document	1-20
A	US 2008/0018881 A1 (HUI et al) 24 January 2008 (24.01.2008) entire document	1-20

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents:	
"A" document defining the general state of the art which is not considered to be of particular relevance	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E" earlier application or patent but published on or after the international filing date	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O" document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

11 January 2018

Date of mailing of the international search report

06 FEB 2018

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
 P.O. Box 1450, Alexandria, VA 22313-1450
 Facsimile No. 571-273-8300

Authorized officer

Blaine R. Copenheaver

PCT Helpdesk: 571-272-4300
 PCT OSP: 571-272-7774