



(51) International Patent Classification:

G06F 17/00 (2006.01) *G06T 7/00* (2006.01)
G06F 3/00 (2006.01) *G06K 9/62* (2006.01)
G06F 3/01 (2006.01) *G06K 9/78* (2006.01)

(21) International Application Number:

PCT/IL2016/050202

(22) International Filing Date:

22 February 2016 (22.02.2016)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/119,226 22 February 2015 (22.02.2015) US

(71) Applicant: **TECHNION RESEARCH & DEVELOPMENT FOUNDATION LIMITED** [IL/IL]; Senate House, Technion City, 3200004 Haifa (IL).

(72) Inventors: **WETZLER, Aaron**; 7/39 Ha-Rav Goren Street, 7575307 Rishon LeTsiyon (IL). **KIMMEL, Ron**; 4/7 Morad Hayasmin Street, 3476202 Haifa (IL).

(81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM,

AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, 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.

(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).

Published:

— with international search report (Art. 21(3))

(54) Title: GESTURE RECOGNITION USING MULTI-SENSORY DATA





FIG. 2

(57) Abstract: A system comprising: a camera configured to capture one or more images of a user's hand; and a computer configured to: receive the one or more captured images, apply a mapping function to the received one or more images, thereby yielding one or more coordinates associated with at least one feature of the user's hand, wherein the mapping function is derived from a set of labeled images that are produced by applying a machine learning algorithm to training data which comprises images of a trainer's hand, wherein the images are labeled with coordinates obtained from multiple magnetic sensors attached to the trainer's hand.

GESTURE RECOGNITION USING MULTI-SENSORY DATA**CROSS-REFERENCE TO RELATED APPLICATION/S**

This application claims the benefit of priority of U.S. Provisional Application
5 No. 62/119,226 filed February 22, 2015 and entitled "Gesture Recognition Using Multi-Sensory Data", the contents of which are incorporated herein by reference in their entirety.

FIELD OF THE INVENTION

10 The invention relates to the field of computerized gesture recognition.

BACKGROUND

Gesture recognition has gained in popularity in a wide range of applications, from user interfaces (UIs) to enhanced computer animation. One common technique to
15 implement gesture recognition relies on supervised machine learning for classifying and identifying gestures performed by a user. However, successful application of this technique requires consideration of several issues.

An adequate training set must be built to produce a reliable classification function. Noise and distortion of the detected gesture and of the underlying training set may
20 inhibit reliable identification and classification.

Hand gestures, with twenty nine degrees of motion are particularly challenging to classify. A machine learning process for such a complex tasks typically requires a large amount of training data where to faithfully represent the camera noise, and hand-finger locations and orientations that would be observed in a non-controlled scenario.

25 Typical techniques to acquire training data use visual markers such as painted gloves, stickers, reflectors or LEDs attached to different parts of the hand or a hand glove for observation by one or more cameras. The positions and orientations of the hand and fingers may then extracted from this visually tracked data. However, visually captured data suffers from occluded markers due to the dimensions and articulation
30 possibilities of the human hand, resulting in incomplete data. Additionally the visual markers themselves may distort the true hand profile and/or impede natural hand motion, compromising the training data.

The foregoing examples of the related art and limitations related therewith are

intended to be illustrative and not exclusive. Other limitations of the related art will become apparent to those of skill in the art upon a reading of the specification and a study of the figures.

5

SUMMARY

The following embodiments and aspects thereof are described and illustrated in conjunction with systems, tools and methods which are meant to be exemplary and illustrative, not limiting in scope.

There is provided, in accordance with an embodiment, a system comprising: a
10 camera configured to capture one or more images of a user's hand; and a computer configured to: receive the one or more captured images, apply a mapping function to the received one or more images, thereby yielding one or more coordinates associated with at least one feature of the user's hand, wherein the mapping function is derived from a set of labeled images that are produced by applying a machine learning
15 algorithm to training data which comprises images of a trainer's hand, wherein the images are labeled with coordinates obtained from multiple magnetic sensors attached to the trainer's hand.

In one embodiment, the feature comprises a digit of the user's hand.

In one embodiment, the feature comprises a joint of the user's hand.

20 In one embodiment, the system further comprises a calibration system that is configured to calibrate the multiple magnetic sensors with a trainer's camera that is configured to capture the images of the trainer's hand.

There is provided, in accordance with an embodiment, a method, comprising:
applying a mapping function to one or more captured images of a user's hand; and
25 producing, from the application of the mapping function, one or more coordinates associated with at least one feature of the user's hand, wherein the mapping function is derived from a set of labeled images that are produced by applying a machine learning algorithm to training data which comprises images of a trainer's hand, and wherein the images are labeled with coordinates obtained from multiple magnetic sensors attached
30 to the trainer's hand.

In one embodiment, the feature comprises a digit of the user's hand.

In one embodiment, the feature comprises a joint of the user's hand.

In one embodiment, the method further comprises calibrating the multiple magnetic sensors with a trainer's camera, and capturing the images of the trainer's hand using the trainer's camera.

There is provided, in accordance with an embodiment, a method for creating a
5 labeled data set for a hand, comprising: obtaining optical image data of a trainer's hand;
obtaining magnetic image data corresponding to the optical image data; identifying a
position attribute from the magnetic image data corresponding to the optical image data,
wherein the attribute is not represented in the optical image data; labeling the optical
image data using the position attribute, thereby creating a labeled image data set; and
10 applying a machine learning algorithm to the labeled image data set to produce a
mapping function.

In one embodiment, the method further comprises calibrating a trainer's camera
with multiple magnetic sensors, wherein the optical image data is captured using the
trainer's camera and wherein the magnetic image data is captured using the multiple
15 magnetic sensors.

In one embodiment, the multiple magnetic sensors comprise low profile sensors.

In one embodiment, the mapping function is applied to an image of a user's hand
to yield one or more coordinates associated with at least one feature of the user's hand.

There is provided, in accordance with an embodiment, a system for creating a
20 labeled data set for a hand, comprising: a camera configured to obtain one or more
optical images of a first user's hand; multiple magnetic sensors that are configured to
attach to the first user's hand and provide magnetic image data corresponding to the
optical image data; and a processor configured to obtain the one or more optical images
and the corresponding magnetic image data to identify a position attribute from the
25 magnetic image data corresponding to the optical image data, wherein the position
attribute is not represented in the optical image data, label the optical image data using
the position attribute, thereby creating a labeled image data set, and apply a machine
learning algorithm to the labeled image data set to produce a mapping function.

In one embodiment, the system further comprises a calibrating system configured
30 to calibrate the camera with the multiple magnetic sensors, wherein the optical image
data is captured using the camera and wherein the magnetic image data is captured
using the multiple magnetic sensors.

In one embodiment, the multiple magnetic sensors comprise low profile sensors.

In one embodiment, the mapping function is applied to an image of a second user's hand to yield one or more coordinates associated with at least one feature of the second user's hand.

5 In one embodiment, the machine learning algorithm comprises one of: a) a convolutional neural network and b) a randomized decision tree.

In one embodiment, a derotation heuristic is applied to the labeled image data set at any of a training phase and a testing phase of the machine learning algorithm.

10 In one embodiment, the derotation heuristic comprises rotating the one or more optical images of the first user's hand such that the base of the thumb of the first user's hand is positioned on the right side of the optical image.

In one embodiment, the processor is configured to label the optical image data by modeling each sensor as a three dimensionally oriented ellipsoid that is raycast into the camera frame, wherein the label for each pixel corresponds to the ellipsoid closest to
15 the camera.

In addition to the exemplary aspects and embodiments described above, further aspects and embodiments will become apparent by reference to the figures and by study of the following detailed description.

20 BRIEF DESCRIPTION OF THE FIGURES

Exemplary embodiments are illustrated in referenced figures. Dimensions of components and features shown in the figures are generally chosen for convenience and clarity of presentation and are not necessarily shown to scale. The figures are listed below.

25 Figs. 1A-1B illustrate images of hands in different articulations;

Fig. 2 illustrates a system for labeling images in accordance with an embodiment;

Fig. 3 illustrates a system for determining multiple position coordinates of an anatomical feature of a hand, in accordance with an embodiment;

30 Fig. 4 illustrates a system for calibrating multiple types of sensory devices, in accordance with an embodiment;

Fig. 5 shows the predicted values vs. ground truth values for the 9 coefficients of the hand orientation matrix;

Fig. 6 shows various orientations for R_{in} , R_{axis} , and α when applying a derotation heuristic;

Fig. 7 illustrates synthetic implementations of applying a derotation heuristic;

Figs. 8A-8B show precision to recall and precision to error threshold for thumb
5 detection using RDT on the HandNet test set; and

Figs. 8C-8D show precision to recall and precision to error threshold for thumb
detection using CNN on the HandNet test set.

DETAILED DESCRIPTION

10 A system and method are disclosed herein to create a set of labeled images that may be applied to a machine learning system for gesture recognition. A set of images, such as may be captured by a standard camera, may be augmented with data obtained via a magnetic sensor apparatus to provide a more robust data set of images that may include information pertaining to otherwise visually obstructed features.

15 The resolution and optical constraints of typical cameras, such as feature obstruction or distortion, as well as noise, pose limitations for reliable classification of gesture recognition. Referring to Figs. 1A-B, two exemplary hand articulations are shown. Fig. 1A shows a hand performing a 'pointing' gesture where the top portions of the digits are obstructed, and Fig. 1B shows a 'wave' gesture, where the back side of the
20 hand is obstructed. Such feature obstruction may cause the omission of important details that impede reliable gesture recognition. Since, the human hand has approximately twenty nine degrees of freedom leading to many possible hand articulations, this poses a limitation for reliable hand gesture recognition.

Reference is now made to Fig. 2 which illustrates a system to combine
25 magnetic imaging with optical imaging to yield a more robust data training set, in accordance with an embodiment. Although the system of Fig. 2 refers to a human hand, any moving target that may be used for gesture recognition. For example, the head, or any limb such as the arm or leg, or digit, or other, may equally be used.

A camera **200** may be provided to obtaining an optical image, such as of a
30 reference hand to create the training set. Additionally, a magnetic tracking system, such as a Trakstar DC magnetic tracking system, may be provided that includes a base transmitter **206** in communication with multiple sensors **202**, such as low profile 2mm

sensors magnetic sensors **202a**, **202b**, and **202c**, that may be affixed to a different joint of the reference hand. DC magnetic trackers such as the TrakStar, that are robust to metallic interference and obstruction by non-ferrous metals, may provide sub-millimeter and sub-degree accuracy for location and orientation relative to a fixed based station. By using low profile sensors, feature distortion may be reduced. Although Fig. 2 shows each finger of the reference hand disposed with two sensors, each corresponding to a different joint, in some embodiments, each finger may be provided with zero, one, or three or more sensors. The palm of the hand may be provided with one or more sensors, such as sensor **202c**. Additional sensors may be affixed to the hand to correspond to additional joints.

Sensors **202** may be affixed to the reference hand via one or more attachers **204**, such as thin elastic loops each provided with a pocket that secures one of sensors **200** to restrict any rotation or change of position relative to attachers **204**. Attachers **204** may each fit on one of the reference hand's fingers in a manner to prevent lateral or medial movement along the finger, and to minimally impact the finger's profile, and not impede the natural motion of the fingers. For example, attachers **204** may fit as 'rings' worn on the fingers of the reference hand. Attaching the sensors using skin tight elastic loops may provide a benefit over gloves in that the depth profile and hand movements are not affected by the attached sensors and thus do not pollute the data obtained by the sensors.

Sensors **202** may provide their orientation with respect to a base transmitter **206** that is provided with the tracking system, and thereby obtaining magnetic image data corresponding to the optical image data. In one embodiment, the transmitter may include a set of three coils that produce electromagnetic fields in three orthogonal directions. The magnetic field produced may be a substantially direct current (DC) magnetic field that may have little or no impact on nearby electronics. A processor **208** may be provided to analyze and process both the magnetic image data and the optical image data to create a labeled image set, as follows: Processor **208** may identify an attribute, such as one or more position coordinates obtained from one of sensors **202** that correspond to a feature of the optical image data. For example, the feature may correspond to a portion of the hand that is obstructed and is therefore not present in the

optical image data. The attribute may be used to label the visual feature, to create the labeled image data set.

In this manner, data obtained from magnetic imaging may provide otherwise missed information regarding features that were obstructed optically. Similarly, the optical image data may provide a visual context to the magnetic image, and may be applied to correct for distortions and/or noise. By combining both the magnetic image data with the optic image data to create an augmented data set, reliability for gesture recognition may be improved.

Multiple images of the hand in different articulations may be obtained and labeled as described above. These labeled images may be collected to create a training set. The training set may be created in a controlled setting such as a lab, where the reference hand of the trainer may be configured with sensors as described above to derive the labeled images included in the training set. Thus created, the training set may be applied to a machine learning algorithm to derive a mapping function. Such a mapping function may be provided to any computer configured with a camera and may be used to determine one or more coordinates of an anatomical hand feature corresponding to a captured image of a user's hand in a non-lab setting, and which may be applied for gesture recognition.

Reference is made to Fig. 3 which illustrates a system for determining multiple position coordinates of an anatomical hand feature, in accordance with an embodiment. A camera **300** configured to capture one or more images of a user's hand **316** may be provided with a computer **308**. Computer **308** may be provided with the mapping function derived by applying a machine learning algorithm to a training data set that includes images labeled with coordinates obtained from multiple magnetic sensors attached to the trainer's hand, as described above. Computer **308** may receive the one or more captured images of the user's hand **316**, and apply the mapping function to the received one or more images, to yield one or more coordinates associated with the user's hand. For example, computer **308** may yield one or more sets of coordinates for the digits or joints of hand **316**. A more detailed implementation of this is given below.

In one embodiment, sensors **202** and base transmitter **206** described in Fig. 2 may be calibrated with camera **200** to calibrate the magnetic image data with the corresponding optical image data and reduce errors and distortion, as follows:

Reference is now made to Fig. 4, which illustrates a system to calibrate multiple types of sensory devices, in accordance with an embodiment. Processor **208** may be provided to calibrate a trainer's camera **200** with a magnetic tracking system comprising multiple sensors **202** and transmitter **206**. Once calibrated, trainer's camera **200**, sensors **202** and transmitter **206** may be used create the training set described above with a lower error and/or distortion. For example, trainer's camera **200** calibrated with sensors **202** may be applied to the system described for labeling multiple images of the trainer's hand to create the training data set.

In one embodiment, the magnetic tracking system may comprise a Trakstar DC magnetic tracking system that includes multiple low profile 2mm magnetic sensors **202** in communication with a transmitter **206**.

Camera **200** may be positioned adjacent to transmitter **206** and may be secured such that camera **200** and transmitter do not move relative to each other. In one embodiment, camera **200** may be mounted on top of, and fastened to transmitter **206**, thereby ensuring that a first coordinate system associated with transmitter **206** is substantially aligned with a second coordinate system associated with camera **200**.

A calibration pattern **210**, such as a checkerboard, may be secured to a board **212**. Board **212** may be of non-ferrous material, such as wood, that does not interfere with a magnetic field. Calibration pattern **210** may include multiple predetermined 'corner' locations **214**, which may be highly visible in both infrared, as well as red, green and blue ('RGB') sensitive cameras and that may be detected with subpixel accuracy. The locations of each of corners **214** relative to camera **200** may be determined during a camera calibration process. Sensors **202** may be affixed to calibration pattern **210** at corners **214** by fitting each of sensors **202** into a hole drilled into board **212** at each of corners **214**. For example, five holes with a diameter of 2mm may be drilled at five corner points **214** on calibration pattern **210**, and each one of sensors **202** may be inserted into each hole. For example, corner points **214** may be situated at the border between the white and black 'checkers' of calibration board **210**. Sensors **202** may be positioned to align the center of each sensor with the surface of calibration board **210**.

and may fit be positioned at the corners **214**. Calibration board **210** may be made of a porous material that 'gives', such as wood or any other suitable material that tightly grips sensors **202** without damaging them, while positioned within the holes. The location of each of corners **214** relative to camera **200** may be determined in a camera calibration step in accordance with conventional methods.

Sensors **202** may be sufficiently small to not extend beyond the surface of calibration board **210**, and not obstruct corner locations **214** from being captured by camera **200**. Sensors **202** may provide their orientation with respect to transmitter **206**, and thus may provide their location and orientation substantially relative to camera **200** positioned adjacent to transmitter **206**.

The above described apparatus may be used to derive a transformation between a coordinate system of camera **200** and a coordinate system of transmitter **206**, thereby enabling converting data obtained from either sensory device to the same frame of reference. For example, the coordinates obtained by transmitter **206** from sensors **202** may be transformed to the frame of reference of camera **200**. The transformed coordinates may be used to annotate, or label a corresponding image captured by camera **200**.

Additionally, the system describe herein may be used to account for any warping resulting from magnetic offset of sensors **402**, or lens distortion from camera **200**. Calibration board **212** may be positioned in multiple calibration poses, such as rotations and/or translations, and may be confined within a predefined distance from camera **200** and transmitter **206** to ensure a substantially uniform deterioration of signal reception by each of camera **200** and transmitter **206** with respect to distance. For each calibration pose, one or more optical images and corresponding magnetic images may be obtained and cross-referenced, as follows:

For each pose, each of sensors **202** may provide a processor **208** a sensor ID in association with a set of coordinates relative to transmitter **210**, corresponding to its subpixel location at each one of corners **214**. Similarly, camera **200** may provide processor **208** a corresponding image of calibration board **210**. Processor **208** may determine a transformation, $M = [R \ t]$, such as a single rigid transform, that converts the location coordinates obtained from sensors **202** and the associated sensor ID to the coordinate system of camera **200** using the location of corners **314** that were determined

during the camera calibration process. In one embodiment, the transformation M may be determined by applying a least squares optimization to the data obtained from sensors **202**.

A non-random offset determined between the coordinate system of camera **200** and the detected positions of sensors **202** may result from a non-linear magnetic distortion in the space in front of camera **200** and/or sensors **202**. This can be corrected by measuring multiple locations for each of sensors **202** and applying a polynomial model to the obtained coordinates to directly compensate for the distortion, such as when correcting lens distortion.

The above steps may be performed by processor **208** to produce a rigid transform between the coordinate system of camera **200** and the coordinate system of sensors **202**. This information may be used to project the locations of each of sensors **202** onto the camera image, and may be used to annotate, or label a region or set of pixels associated with the sensor ID.

In another embodiment, a common clock such as may be integrated with processor **208** may be defined to correctly associate information obtained from sensors **202** with the images obtained from camera **200**. This may be especially relevant to moving objects whose optical and magnetic images change with time.

In another embodiment, the coordinates obtained from sensors **202** and the image obtained from camera **200** may be reverse matched by maintaining a buffer to store the most recent signals from both sensors **202** and camera **200**. These signals may be matched in real time, but matching synchronously obtained signals.

For example, the Trakstar system can sample data at 420 Hz which is considerably faster than then camera frame rate of 60Hz. Thus, a buffer of 7 Trakstar samples may be maintained for each of sensors **202**. The most recent camera image may be matched to every seventh most recently obtained sample from sensors **202**. These matched samples may be transformed as described above, and the image may be labeled with the synchronized magnetic signal.

In one embodiment, the 3D location coordinates of corners **214** points may be determined relative to the coordinate system of camera **200** using standard calibration techniques. These same 3D location coordinates for corners **214** may be known relative to the coordinate system of transmitter **206**, *a priori*, as a result

of their intentional positioning at predefined corners **214**, resulting in a set of matched location coordinates in 3D for both the coordinate system of camera **200** as well as the coordinate system for transmitter **206**. These matched location coordinates may be applied to derive the transformation between the camera system and the magnetic sensors and transmitter system.

In an embodiment, the labeled images may be collected in a database that may be applied to extract highly accurate hand and finger positions and orientations as well as pixel labels for each individual fingertip and hand and wrist joints. Such a database may be applied in a machine learning system for gesture recognition, where Deep Convolutional Networks and a Randomized decision trees may be taught to predict an orientation of the hand and location of the fingertips in three dimensions (3D) when coupled with a three-dimensional camera.

For example, the labeled images obtained using the DC magnetic tracking system, such as the Trakstar system as described above, may be used to build a large-scale database of annotated depth images of different hand orientations having three degrees of freedom (DOF).

The publication '*Rule of Thumb: Deep Derotation for improved fingertip detection*', Aaron Wetzler et al, *Proceedings of the British Machine Vision Conference (BMVA Press)*, pages 33.1-33.12, Sept. 2015, (<https://dx.doi.org/10.5244/C.29.33>) is hereby incorporated by reference in its entirety.

In one exemplary implementation using an Intel RealSense camera with 2mm TrakStar magnetic sensors, the method and system described above was used to create the HandNet database, available as of February 22, 2016 at <http://www.cs.technion.ac.il/~twerd/HandNet/>, and having 212,928 unique hand poses. The sensors were attached to each user's hand and fingertips with tight elastic loops and contained within sewn seam pockets to prevent lateral and medial movement of the sensors along the finger. Since the skin tight elastic loops did not affect the depth profile of the hand, or hinder the movement of the hand, they did not pollute the data as would gloves.

To create such a database, the camera and sensor frame may be calibrated as described above, by positioning the magnetic sensors on the corners of a checkerboard pattern to create a physical correspondence between the detected corner locations and

the actual sensors. The extracted 2D locations of the corner points on the calibration board, together with the sampled sensor 3D locations, may be used to solve the Perspective-n-Point problem (PnP) using the technique described in *V. Lepetit, F. Moreno-Noguer, and P. Fua. E-pnp: An accurate $O(n)$ solution to the pnp problem. International journal of computer vision, pages 155–166, 2009*, to determine the extrinsic configuration between the devices.

To annotate, or label, each image, each sensor may be modeled as a 3D oriented ellipsoid. The ellipsoid may be raycast into the camera frame and the annotation, or label for every pixel, may be set to correspond to the identity of the ellipsoid (sensor) closest to the camera. A heatmap h_i may be created for each fingertip i using a similar technique, with the noted difference that the value per pixel may be defined as a gaussian distribution over the distance to the projected sensor location.

The annotated database may be created from images taken of multiple different, diverse users. For example, the HandNet database was created from images obtained from ten participants wearing the magnetic sensors: five males and five females, all having varying hand sizes. Images were captured while the participants performed random hand motions with extensive pose variation. The RealSense camera operated at a 58 frames per second (fps) rate, and produced 640×480 pixel depth maps, which were reduced to 320×240 pixels. The TrakStar system sampled measurements at a rate of 720Hz to record a total of 256,987 images, of which a portion were removed due to low quality, yielding a final dataset of 212,928 frontal hand images having annotations for the position and orientation for each fingertip as well as the back of the palm. For each participant, the rotation and location of each sensor was adjusted with an offset for positional consistency.

Such a database may be used to train computer learning systems, such as a Deep Convolutional Neural Network (CNN) or a Randomized Decision Tree (RDT), which may be applied in a gesture recognition application to predict the orientation and location of the fingertips of a hand recorded by a 3D camera. Since hand gesture recognition poses challenges due to a) the high propensity for occlusion, and b) the many rotations that vary the appearance of the object parts, the hand space variability during both the training phase and run-time of the gesture recognition application may be substantially reduced. For example, the rotations may be learned in a regression

context using the CNN, such as described in *J. Tompson, M. Stein, Y. Lecun, and K. Perlin. Real-time continuous pose recovery of human hands using convolutional networks. ACM Transactions on Graphics (TOG), 33, 2014.*

To this end, a method for normalizing out the effects of rotation on highly articulated motion of deforming geometric surfaces, such as hands observed by a depth camera, is now described. Predicted orientation is combined with an in-plane derotation scheme that utilizes an inherent characteristic of hand rotations, namely that there is almost always an in-plane rotation which can be applied to an image of the hand which results in the positioning of the base of the thumb to be on the right side of the image.

Derotation in this context may be understood to be the process of warping an input image of a hand by the inverse of the detected inplane rotation of the projection of the hand onto the image. The rotation is measured according to a canonical reference rotation and therefore the derotation process attempts to rotate the projected image back to its canonical rotation. The canonical frame of the hand in an image may be roughly defined as a hand orientation that positions the thumb on the right side of the image. The observation that an image of a hand can almost always be rotated so that the thumb is seen on the right is the insight which enables the "Rule of thumb" derotation procedure.

Derotating the hand image to such a canonical pose, instead of augmenting a dataset with all variations of the rotational degrees of freedom, may allow overcoming the ambiguity inherent in rotationally variant features.

This derotation scheme can be used to predict three degrees of freedom (DOF) orientation information on a database of hand images captured by a depth sensor, allowing location and identification of the fingertips. The detections may be overlaid on IR images from the camera which is not part of the classification process.

In one embodiment, instead of a creating a heatmap for determining positioning, the 9 coefficients of a rotation matrix may be directly predicted. A regular rotation is characterized by 3 degrees of freedom. However, by using 9 parameters with a large database, the over-parameterized output may be effectively regularized. This representation of the rotation matrix is unique in the group $SO(3)$, unlike quaternions and Euler angles which may be noisy and unreliable, particularly when used to predict a single representative angle. For training purposes, Euclidian loss may be used and

orthonormality not enforced. The output rotation matrix R of this CNN may be directly projected onto the closest unitary matrix using singular value decomposition (SVD), to yield the orientation prediction matrix $\hat{R} = [r_1, r_2, r_3]$, where $R = USV^T$. Enforcing the determinant, $\det(\hat{R}) = 1$ yields $\hat{R} = UV^T$ as a least squares optimal projection into the group $SO(3)$.

Reference is made to Fig. 5 which shows the result of predicting the 9 ground truth coefficients for the HandNet database relative to the ground truth. For clarity each ground truth coefficient was ordered monotonically. This reordering was applied to the predicted results. The mean squared error for all the coefficients on the HandNet test set before and after SVD is 0.0271 and 0.0234 respectively

To reduce pose variance, a heuristic may be applied to rotate the image of the hand such that the thumb is positioned on the right side of the image. To accomplish this, the camera image may be rotated about its center by an angle α computed using the orientation prediction $\hat{R} = [r_1 r_2 r_3]$.

Referring to Fig. 6, the angle α with which to rotate the image may be set as the angle between the projection of a predefined axis, defined by R_{axis} and the upwards image direction, defined by R_{in} . However, when this axis points to or away from the camera the projection onto the screen will be small and noisy. Thus, a derotation heuristic that selects an alternative axis when this condition holds may be applied. Specifically, the predicted axis r_{align} most aligned with the camera z axis may be determined as $r_{align} = \operatorname{argmax}_{r_i \in \{r_1, r_2, r_3\}} \|(0,0,1) \cdot r_i\|$. If r_{align} is either in the palm pointing direction or the in direction of the extended fingers, then r_2 pointing in the thumb direction may be non-noisy, and r_{orient} may be set as $r_{orient} = r_2$. However, if $r_{align} = r_2$, such as if the thumb direction is mostly pointing towards or away from the camera, then r_{orient} may be set as $r_{orient} = r_3$, corresponding to the palm vector. An implementation of this heuristic is illustrated procedure in the following Derotation Algorithm:

Procedure Derotate (R)

$$r_{align} \leftarrow \operatorname{argmax}_{r_i \in \{r_1, r_2, r_3\}} \|(0,0,1) \cdot r_i\|$$

If $r_{align} = r_2$ (thumb aligned axis) then

$$\alpha \leftarrow \operatorname{atan2}(r_{3x}, r_{3y}) + 90 + \begin{cases} 180 & \text{if } r_{2z} \leq 0 \\ 0 & \text{if } r_{2z} > 0 \end{cases}$$

Else

$$\alpha \leftarrow \text{atan2}(r_{2x}, r_{2y}) + 90$$

End if

Return α

5 **End Procedure**

Synthetic examples of applying this heuristic are shown in Fig. 7. It may be noted that this technique is adaptable for objects other than the hand. DeROT may be defined as a method that combines using a convolution neural network to predict the orientation $\hat{R}$, while applying this derotation heuristic. Fig. 7 shows the depth projection (700, 702, 704) of the virtual hand before applying DeROT which can be seen on the left wall of the cube representing the camera plane. The axis marked r_{orient} is projected onto the camera plane and used in DeROT to define the angle α . The shaded circles (706, 708, 710) contain the resulting images of the hands after applying derotation by angle α , resulting in the thumb's positioning on the right side of the image.

15 Optionally to using the DeROT heuristic, the principal axes of the hand silhouette may be extracted using Principal Component Analysis (PCA), and the rotation angle of the largest axis to the vertical image axis may be determined as α . The image may then be rotated by angle α about this largest axis accordingly.

Another option may be to perform a segmentation on the hand image, described in greater detail below, and determine an enclosing ellipse, such as by using a Procrustes-type algorithm over the convex hull of the points V of the hand segmentation. The minimum area of the enclosing ellipse may be found efficiently over the points $x_i \in \text{convhull}(V)$ by minimizing $-\log(\det(A))$, s. t. $(x_i - \bar{x}_i)^T A (x_i - \bar{x}_i) \leq 1$ for $A, \bar{x}_i$ defining the ellipse. This may be solved using Khachiyan's algorithm, as described in *B. Aspvall and R. Stone. Khachiyan's linear programming algorithm. pages 1–13, 1980.*

A method to perform fingertip detection on a per frame basis for each captured depth image, without requiring frame-to-frame tracking or kinetic modelling is now described. The target hand may be segmented from the depth image using either a fast depth-based flood-fill method seeded either from the previous frame for real-time use and testing, or from the ground truth hand location for building the database. Using the center of mass (CoM) of the segmented hand and its average depth value, a depth

dependent bounding box may be defined corresponding to the camera used to capture the depth image. For example, a bounding box of size $w = \frac{50000}{z}$ may be defined for a RealSense camera (HandNet) and $w = \frac{70000}{z}$ for a Kinect camera (NYUHands) where z is the depth of the CoM of the segmented hand. The image may be derotated about the CoM using an angle of rotation α determined as described above, such as by determining α according to the in-plane angle calculated by applying the DeROT heuristic above. This comes from the predicted full 3D orientation at run-time or from the ground truth sensor orientation for database construction or testing. The image may then be cropped using the bounding box. A learning-based fingertip detector may be used to detect the fingertips on the cropped image, accordingly, such as a random decision tree, or convolutional neural network.

When using a Random Decision Tree learning-based detector to detect the fingertips, methods may be used such as described in *C. Keskin, F. Kiraç, Y. Emre Kara, and L. Akarun, "Real time hand pose estimation using depth sensors", in Consumer Depth Cameras for Computer Vision, pages 119–137 Springer, 2013*, and *J. Shotton, A. Fitzgibbon, M. Cook, T. Sharp, M. Finocchio, R. Moore, A. Kipman, and A. Blake, "Real-time human pose recognition in parts from single depth images", Conference on Computer Vision and Pattern Recognition (CVPR), pages 297–1304, 2011*.

Optionally, the random binary depth attributes per pixel may be spatially distributed according to an exponential sampling pattern, such as that described by *L. Stefan, C. Margarita, and S. Roland Yves in "Brisk: Binary robust invariant scalable keypoints", International Conference on Computer Vision (ICCV), pages 2548–2555, IEEE, 2011*. Additionally, in one embodiment, only a single RDT may be used. After training the single RDT, the class distributions representing the empirical estimate of the posterior probability $p(c|x)$ of hand part label c given the image evidence x stored at each leaf may be used for inference. Inferring the most likely fingertip identity label may be performed pixel-wise by finding the label c^* which maximizes $p(c|x)$ per pixel. To address noise, such as may result from a lack of influence of neighboring classifications, a spatial regularization may be applied as follows:

- for each fingertip i , treat the posterior $p(c=i|x)$ for all pixels of the fingertip as an image, and convolve the image with a discrete two –dimensional (2D)

gaussian smoothing kernel g_σ having blur radius σ . This may correlate the posterior label distributions of nearby pixels. Therefore every pixel q may be labeled by fingertip identity (including palm and wrist labels) according to

$$c^*(q, x) = \operatorname{argmax}_{i \in \{0..6\}} (g_\sigma * p_{c=i|x})(q) \quad (1)$$

5 The close proximity of fingers may compromise standard mean-shift clustering. Thus, the largest label regions, or ‘blobs’ in the label image from Equation 1 may be detected above a certain area threshold. The 2D fingertip locations may then be assigned to the blob centers and, if necessary, the average depth value for each blob may be used to generate the 3D camera-space coordinates.

10 Training optimal decision trees is known to be NP-complete and therefore trees are built from the root down using breadth-first greedy optimization over tree node impurity. Optionally, the Gini impurity measure which is slightly cheaper to compute than the more typical entropy measure may be applied. To build a database for training an RDT, a predefined portion of the fingertip pixels and non-fingertip hand pixels may
15 be extracted. For example, to build the HandNet database, 80% of the fingertip pixels and 50% of the non-fingertip hand pixels may be extracted from the training datasets, resulting in a training dataset of 500 million sample pixels totaling 600GB of data for 1200 attributes. In the HandNet example, a tree-building algorithm trained an unpruned randomized tree on $4 \times$ GTX 580 graphic processing units (GPUs) and an Intel I7
20 processor with 48GB of RAM in 16 hours for a tree depth of 21 with 18000 query tests per node. It may be note that a very large number of samples may help to prevent overfitting of single RDTs.

25 When using a CNN architecture, methods such as described in *J. Tompson, M. Stein, Y. Lecun, and K. Perlin, 'Real-time continuous pose recovery of human hands using convolutional networks', ACM Transactions on Graphics (TOC), 33, 2014*, and supplementary material may be used to predict the location of the five fingertips by using the maximum location in a set of heat maps that implicitly represent fingertip locations. Using the deeper convolutional layers provided by the CNN may allow extracting feature responses of highly articulated objects, such as oriented fingertips, at
30 a higher semantic level. The heatmap based error objective may help to spatially regularize the network during training. For input to the CNN, D_1 may set to be the

cropped depth of the image, resized to 96×96 pixels. D_1 may be downsample twice, each time by a factor of two, to produce D_2 and D_3 , respectively. A subtractive form of local contrast normalization (LCN), such as described in *K. Jarrett, K. Kavukcuoglu, M. Ranzato, and Y. LeCun, 'What is the best multi-stage architecture for object recognition?', International Conference on Computer Vision (ICCV), pages 2146–2153. IEEE, 2009* and *Tompson (above)*, may be used such that $D_i \leftarrow D_i - g_\sigma * D_i$ using a gaussian smoothing kernel g_σ with $\sigma = 5$ pixels. The triplet (D_1, D_2, D_3) may then be the input to the neural network. The trained network may output a heatmap h_i per fingertip i for new data.

Optionally, the output for a non-fingertip heatmap that is strong may be augmented wherever a fingertip is not likely to be present. Additionally, instead of fitting a gaussian model to the strongest mode in the low resolution heatmaps, each 18×18 pixel fingertip heatmap h_i may be upsampled to a fixed size of 128×128 pixels using a smoothing bi-linear interpolator. Similar to the technique described above, every pixel q may be labeled with fingertip identity (including a non-fingertip class)

$$c^*(q) = \underset{i \in \{0..5\}}{\operatorname{argmax}} h_i(q) \quad (2)$$

where the fingertip locations are given by the location of the largest label blob.

To train the CNN, both the orientation regression CNN described above as well as the described fingertip CNN may be trained using any suitable technique, such as described in *Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, 'Caffe: Convolutional architecture for fast feature embedding', arXiv preprint arXiv:1408.5093, 2014* using an NVidia GTX 980 with an i7 processor and 16GB of onboard RAM. Training may be done both with a Euclidean loss and a batch size of 100 for 100000 iterations with stochastic gradient descent. The initial learning rate may be in the order of 0.01, and may be reduced every predetermined number of iterations, such as by a factor of 0.2 for every ten thousand iterations. Such repeated fine-tuning may help the network converge.

Experimental Results

Experiments to test the method and system described above were performed separately using two databases: the HandNet database and the publicly available

database NYUHands. The baseline results (1) were obtained by training on non-derotated data and testing on non-derotated data. This baseline was compared to results obtained by training on non-derotated data while testing with derotated data (II), training on derotated data while testing with non-derotated data (III), and training on
5 derotated data while testing with derotated data (IV).

For the non-derotated data, 202,928 images were randomly selected from the HandNet database for training, and the remaining 10,000 images were used for testing. For the NYUHands database, all three camera views (72,757 images per view) were used for training purposes, and the frontal view was used for testing (8,252 images).
10 The training and testing sets were slightly diluted in accordance with the hand segmentation technique described above, resulting in 184,100 training images and 7,241 testing images.

Non-derotated data was used to train two CNN orientation regression networks, one for each dataset for experiment types (2) and (4), as well as for training the RDT
15 and CNN fingertip detectors for experiment types (1) and (2). However, for testing the fingertip detectors in experiments (1) and (3) (testing on non-derotated data), each testing image was rotated by uniformly random in-plane rotational offsets ranging between -90 and 90 degrees, to guarantee that the testing data differs from the training data.

Derotated data was used for the training data in experiment types (3) and (4) (training on derotated data). The data was first derotated by an Oracle, implemented by DeROT defined above using the ground truth R_{gt} obtained from the magnetic sensors. With experiment types (2) and (4) (testing on derotated data), the same uniform random image rotation was first applied to the test images as for experiment types (1) and (2).
25 One of the following derotation techniques was then applied: (a) Procrustes derotation, (b) DeROT using $\hat{R}$ predicted by the CNN regression network, (c) Oracle derotation with R_{gt} .

Mean precision and mean average precision: Precision and recall were computed according to the protocol described in *M. Everingham, L. Van Gool, J. Winn, C. Williams, and A. Zisserman, 'The PASCAL Visual Object Classes Challenge 2011 (VOC2011) Results', (http://www.pascalnetwork.org/challenges/VOC/voc2011/workshop/index.html)*. The prediction confidence was set as the value at the location of
30

the fingertip detection in the 128×128 channel heatmap for each fingertip. The mean precision (mP) represents the mean precision over all fingertips at a recall rate of 100%. Mean average precision (mAP) measures the mean of all the areas under the precision-recall curves for each fingertip and takes into account the behaviour over all confidence values.

Error threshold: The error of a prediction was defined as the distance to the ground truth location. False positives were considered to be when the predicted distance of a fingertip was more than 6 pixels from the ground truth position. The threshold of 6 pixels roughly translates into a distance of 1cm for both the HandNet and NYUHands databases in an image patch of size 128×128 pixels cropped as described above. 1cm was determined to be a natural threshold, accounting for an average distance between adjacent fingertips of over 1.6cm on average.

The following table (Table 1) summarizes the experimental results for the experiments described above:

Test set Derotation Method	None		(a) Procrustes		(b) DeROT		(c) Oracle	
	mP	mAP	mP	mAP	mP	mAP	mP	mAP
HandNet								
RDT trained on non-derotated data	0.51	0.79	0.49	0.77	0.55	0.85	0.60	0.87
RDT trained on derotated data	0.32	0.60			<u>0.63</u>	<u>0.88</u>	<u>0.75</u>	<u>0.95</u>
CNN trained on non-derotated data	0.44	0.73	0.42	0.73	0.46	0.77	0.50	0.79
CNN trained on derotated data	0.30	0.59			<u>0.61</u>	<u>0.88</u>	<u>0.74</u>	<u>0.95</u>
NYUHands								
RDT trained on non-derotated data	0.51	0.75	0.47	0.73	0.58	0.84	0.61	0.86
RDT trained on derotated data	0.35	0.58			<u>0.63</u>	<u>0.88</u>	<u>0.68</u>	<u>0.89</u>
CNN trained on non-	0.38	0.70	0.36	0.69	0.46	<u>0.80</u>	0.48	<u>0.81</u>

derotated data								
CNN trained on derotated data	0.23	0.42			<u>0.49</u>	0.72	<u>0.53</u>	0.73

Table 1

The results in bold outperformed the baseline values shown in italics. For each row pair (derotated training data vs non-derotated training), the underlined result indicates the better of the two. For each row pair (derotated training data vs non-derotated training), the underlined result is the better of the two.

From the results shown in Table 1, Procrustes consistently reduces the quality of fingertip detection. Conversely, DeROT outperforms the baseline for every experiment. For all but one experiment, this improved performance is significantly enhanced by training on derotated data instead of original data. The results from the Oracle serve as an upper bound achievable by derotation.

Furthermore, the experiments indicates that the deep derotation method (DeROT) gave an improvement of 20.5% in mean average precision (mAP) over the baseline results for two fingertip detection in depth image approaches: a random decision tree (RDT) and a deep convolutional neural network (CNN). These results were compared to a non-learning based method similar to PCA, and which produced inferior results, indicating the improvement in fingertip detection using the DeROT method described above.

Reference is now made to Figs. 8A-8D, which show precision to recall and precision to error threshold for thumb detection using RDT and CNN on the HandNet test set. Each line indicates an experiment which is labeled in the legend using the experiment types outlined above, and the derotation types Procrustes(a), DeROT(b), Oracle(c). Line **800** indicates the baseline. Training on derotated data and then applying DeROT (line **802**), or Oracle (line **804**) are indicated accordingly. Training on non-derotated data and then applying DeROT (line **806**) or Oracle (line **808**) are indicated accordingly. The average precision (AP) and precision at 1cm error (P@1cm) are shown for each thumb experiment.

These results show that the use of DeROT improves over the baseline results for all measurements for both RDT and CNN for experiments on both datasets. On

HandNet, when training an RDT and CNN on ground truth derotated data, test-time use of DeROT yields improvement in mAP of 11.3% and 20.5% over the respective baselines. For NYUHands, DeROT gives an RDT a gain of 17.3% in mAP when trained on derotated data and a CNN achieves mAP gains of 14.2% when trained on underrotated data but only a marginal gain of 2.5% when trained on derotated data. The confidence values for this specific case were not found to be reliable (which directly effects mAP) because of confusion between fingertips (specifically index and ring) and which further justified the creation of the HandNet database. For all experiments and datasets, using DeROT improves the mP results between 7.8% and 21.1% on underrotated training data, and between 23.5% and 38.6% for derotated training data. For these experiments a single RDT mostly outperforms a CNN, albeit they are trained with different data and objectives.

Thus, using derotation, specifically DeROT, may significantly improve the localization ability of machine-learning based per-frame fingertip detectors by reducing the variance of the pose space. Furthermore this procedure may work despite the extremely high range of potential poses. This approach may provide an alternative to data augmentation and as a potentially useful additional step in pipelines dedicated to articulated object pose extraction such as hands.

Additionally, the system and method described above may be relevant to mobile devices with 3D cameras for calibrating and learning by different devices in different lighting scenarios and skin tones and/or 3D sensor characteristics. One advantage of the system disclosed herein is the direct use of exact camera data in a real recording scenario, thereby eliminating the need for less accurate or less reliable simulated data. Thus, the annotated data obtained by the methods described above may account for the noise and disturbance characteristics of real camera images in real situations, which may be difficult to simulate accurately.

Furthermore, the system and method described above may be applied for recording egocentric motion in a virtual reality or augmented reality scenario. In particular, these systems and methods may be used for training such a virtual and/or augmented reality system. The true motion of a trainer's hands may be recorded from a first-person point-of-view by attaching the camera to the trainer's head such that the camera faces the trainer's hands. Using a depth generating camera system, such as a

structured light scanner or stereo camera reconstruction setup to obtain such automatically labeled data for use by a machine learning system may provide a powerful tool for use by Virtual Reality (VR) and AR systems where markerless detection and tracking of a user's hands may be critical for a useful user experience.

5 The present invention may be a system, a method, and/or a computer program product. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

10 The computer readable storage medium can be a non-transitory, tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the
15 computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically
20 encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media
25 (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the
30 Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A

network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

5 Computer readable program instructions for carrying out operations of the present invention may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, or either source code or object code written in any combination of one or more programming languages, including an object oriented
10 programming language such as Java, Smalltalk, C++ or the like, and conventional procedural programming languages, such as the "C" programming language or similar programming languages. The computer readable program instructions may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on
15 the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example,
20 programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present invention.

25 Aspects of the present invention may be described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be
30 implemented by computer readable program instructions.

These computer readable program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data

processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also
5 be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram
10 block or blocks.

The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the
15 instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and
20 computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the block may occur out of the order noted in
25 the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose
30 hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

The descriptions of the various embodiments of the present invention have been

- presented for purposes of illustration, but are not intended to be exhaustive or limited to the embodiments disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the described embodiments. The terminology used herein was chosen to best explain the
- 5 principles of the embodiments, the practical application or technical improvement over technologies found in the marketplace, or to enable others of ordinary skill in the art to understand the embodiments disclosed herein.

CLAIMS

What is claimed is:

1. A method for creating a labeled data set for a hand, comprising:
 - 5 obtaining optical image data of a trainer's hand;
 - obtaining magnetic image data corresponding to the optical image data;
 - identifying a position attribute from the magnetic image data corresponding to the optical image data, wherein the attribute is not represented in the optical image data;
 - 10 labeling the optical image data using the position attribute, thereby creating a labeled image data set; and
 - applying a machine learning algorithm to the labeled image data set to produce a mapping function.
- 15 2. The method of claim 1, further comprising calibrating a trainer's camera with multiple magnetic sensors, wherein the optical image data is captured using the trainer's camera and wherein the magnetic image data is captured using the multiple magnetic sensors.
3. The method of claim 2, wherein the multiple magnetic sensors comprise low
20 profile sensors.
4. The method of claim 1, wherein the mapping function is applied to an image of a user's hand to yield one or more coordinates associated with at least one feature of the user's hand.
5. A system for creating a labeled data set for a hand, comprising:
 - 25 a camera configured to capture one or more optical images of a first user's hand;
 - multiple magnetic sensors that are configured to attach to the first user's hand and provide magnetic image data corresponding to the optical image data; and
 - 30 a processor configured to:

obtain the one or more optical images and the corresponding magnetic image data to identify a position attribute from the magnetic image data corresponding to the optical image data, wherein the position attribute is not represented in the optical image data,

5 label the optical image data using the position attribute, thereby creating a labeled image data set, and

 apply a machine learning algorithm to the labeled image data set to produce a mapping function.

10 6. The system of claim 5, further comprising a calibrating system configured to calibrate the camera with the multiple magnetic sensors, wherein the optical image data is captured using the camera and wherein the magnetic image data is captured using the multiple magnetic sensors.

7. The system of claim 6, wherein the multiple magnetic sensors comprise low profile sensors.

15 8. The system of claim 5, wherein the mapping function is applied to an image of a second user's hand to yield one or more coordinates associated with at least one feature of the second user's hand.

9. The system of claim 5, wherein the machine learning algorithm comprises one of: a) a convolutional neural network and b) a randomized decision tree.

20 10. The system of claim 10, wherein a derotation heuristic is applied to the labeled image data set at any of a training phase and a testing phase of the machine learning algorithm.

11. The system of claim 11, wherein the derotation heuristic comprises rotating the one or more optical images of the first user's hand such that the base of the thumb of the first user's hand is positioned on the right side of the optical image.

25 12. The system of claim 5, wherein the processor is configured to label the optical image data by modeling each sensor as a three dimensionally oriented ellipsoid that is raycast into the camera frame, wherein the label for each pixel corresponds to the ellipsoid closest to the camera.

30

13. A system comprising:

a camera configured to capture one or more images of a user's hand; and

a computer configured to:

receive the one or more captured images,

5 apply a mapping function to the received one or more images, thereby
yielding one or more coordinates associated with at least one feature of the
user's hand, wherein the mapping function is derived from a set of labeled
images that are produced by applying a machine learning algorithm to
training data which comprises images of a trainer's hand, wherein the
10 images are labeled with coordinates obtained from multiple magnetic
sensors attached to the trainer's hand.

14. The system of claim 13, wherein the feature comprises a digit of the user's hand.

15. The system of claim 13, wherein the feature comprises a joint of the user's hand.

16. The system of claim 13, further comprising a calibration system that is
15 configured to calibrate the multiple magnetic sensors with a trainer's camera
that is configured to capture the images of the trainer's hand.

17. A method, comprising:

applying a mapping function to one or more captured images of a user's
hand; and

20 producing, from the application of the mapping function, one or more
coordinates associated with at least one feature of the user's hand, wherein
the mapping function is derived from a set of labeled images that are
produced by applying a machine learning algorithm to training data which
comprises images of a trainer's hand, and wherein the images are labeled
25 with coordinates obtained from multiple magnetic sensors attached to the
trainer's hand.

18. The method of claim 17, wherein the feature comprises a digit of the user's hand.

19. The method of claim 17, wherein the feature comprises a joint of the user's hand.

20. The method of claim 17, further comprising calibrating the multiple magnetic sensors with a trainer's camera, and capturing the images of the trainer's hand using the trainer's camera.

1/11




FIG. 1A




FIG. 1B

2/11




FIG. 2

3/11




FIG. 3

4/11




FIG. 4

5/11




FIG.5

6/11




FIG.6




FIG. 7

8/11





FIG. 8A

9/11



RESULT
USING RDT

FIG.8B

10/11





FIG. 8C

11/11



RESULT
USING CNN

FIG.8D

INTERNATIONAL SEARCH REPORT

International application No.

PCT/IL2016/050202

A. CLASSIFICATION OF SUBJECT MATTER IPC (2016.01) G06F 17/00, G06F 3/00, G06F 3/01, G06T 7/00, G06K 9/62, G06K 9/78 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) IPC (2016.01) G06F 17/00, G06F 3/00, G06F 3/01, G06K 9/62, G06K 9/78, G06T 7/00 Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) Databases consulted: Esp@cenet, Google Patents, Google Scholar, FamPat database Search terms used: gesture recognition hand magnet camera sensor marker coordinates map optic correspond combine		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	"Sensors for Gesture Recognition Systems"; IEEE Transactions on Systems Man and Cybernetics Part C (Applications and Reviews); 42(3); pp.277-290; Retrieved from the Internet: <URL: https://www.researchgate.net/profile/Helman_Stern/publication/254058478_Sensors_for_Gesture_Recognition_Systems/links/53f37c3d0cf256ab87b0aff0.pdf > BERMAN Sigal; STERN Helman 03 May 2012 (2012/05/03) The whole document	1-20
A	US 8543240 B2 Itkowitz et al. 24 Sep 2013 (2013/09/24) Abstract; column 3, lines 6-64; column 14, lines 20-55; column 15, lines 12-67	1-20
A	WO 2014058909 A2 OBLONG INDUSTRIES [US] 17 Apr 2014 (2014/04/17) Abstract; page 54, lines 2-29; page 55, line 6 - page 58, line 21	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 "E" earlier application or patent but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "&" document member of the same patent family		
Date of the actual completion of the international search 31 May 2016		Date of mailing of the international search report 01 Jun 2016
Name and mailing address of the ISA: Israel Patent Office Technology Park, Bldg.5, Malcha, Jerusalem, 9695101, Israel Facsimile No. 972-2-5651616		Authorized officer PLACHINTA Ekaterina Telephone No. 972-2-5651740

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/IL2016/050202

Patent document cited search report			Publication date	Patent family member(s)			Publication Date
US	8543240	B2	24 Sep 2013	US	2011118753	A1	19 May 2011
				US	8543240	B2	24 Sep 2013
				BR	112012011277	A2	12 Apr 2016
				BR	112012011321	A2	19 Apr 2016
				BR	112012011324	A2	19 Apr 2016
				BR	112012011422	A2	03 May 2016
				CN	102596085	A	18 Jul 2012
				CN	102596085	B	17 Jun 2015
				CN	102596086	A	18 Jul 2012
				CN	102596086	B	20 Jan 2016
				CN	102647955	A	22 Aug 2012
				CN	102665588	A	12 Sep 2012
				CN	102665588	B	22 Jul 2015
				CN	102665589	A	12 Sep 2012
				CN	102665589	B	30 Sep 2015
				CN	104856764	A	26 Aug 2015
				CN	104887326	A	09 Sep 2015
				EP	2467082	A2	27 Jun 2012
				EP	2467082	B1	27 Apr 2016
				EP	2480156	A1	01 Aug 2012
				EP	2480157	A1	01 Aug 2012
				EP	2480158	A1	01 Aug 2012
				EP	2480158	B1	11 May 2016
				EP	2498711	A2	19 Sep 2012
				JP	2013510675	A	28 Mar 2013
				JP	5699158	B2	08 Apr 2015
				JP	2013510672	A	28 Mar 2013
				JP	5702797	B2	15 Apr 2015
				JP	2013510673	A	28 Mar 2013
				JP	5702798	B2	15 Apr 2015
				JP	2013510671	A	28 Mar 2013

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/IL2016/050202

Patent document cited search report	Publication date	Patent family member(s)	Publication Date
		JP 5757955 B2	05 Aug 2015
		JP 2013510674 A	28 Mar 2013
		JP 5847725 B2	27 Jan 2016
		JP 2015107377 A	11 Jun 2015
		JP 2015128681 A	16 Jul 2015
		JP 2015186651 A	29 Oct 2015
		KR 20120102647 A	18 Sep 2012
		KR 20120115484 A	18 Oct 2012
		KR 20120115487 A	18 Oct 2012
		KR 20130027006 A	14 Mar 2013
		KR 20140015144 A	06 Feb 2014
		US 2011118748 A1	19 May 2011
		US 8521331 B2	27 Aug 2013
		US 2011118752 A1	19 May 2011
		US 8682489 B2	25 Mar 2014
		US 2014018960 A1	16 Jan 2014
		US 8831782 B2	09 Sep 2014
		US 2012071892 A1	22 Mar 2012
		US 8935003 B2	13 Jan 2015
		US 2012071891 A1	22 Mar 2012
		US 8996173 B2	31 Mar 2015
		US 2015080909 A1	19 Mar 2015
		US 2015182289 A1	02 Jul 2015
		WO 2011060139 A2	19 May 2011
		WO 2011060139 A3	19 Apr 2012
		WO 2011060171 A1	19 May 2011
		WO 2011060185 A1	19 May 2011
		WO 2011060187 A1	19 May 2011
		WO 2012044334 A2	05 Apr 2012
		WO 2012044334 A3	31 May 2012

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/IL2016/050202

Patent document cited search report	Publication date	Patent family member(s)	Publication Date
WO 2014058909 A2	17 Apr 2014	WO 2014058909 A2	17 Apr 2014
		WO 2014058909 A3	17 Jul 2014
		CN 101536494 A	16 Sep 2009
		CN 101689152 A	31 Mar 2010
		CN 102047203 A	04 May 2011
		CN 102112945 A	29 Jun 2011
		CN 102177513 A	07 Sep 2011
		CN 102177513 B	02 Dec 2015
		CN 102224476 A	19 Oct 2011
		CN 102460510 A	16 May 2012
		CN 102460510 B	20 Apr 2016
		CN 102804206 A	28 Nov 2012
		CN 103930944 A	16 Jul 2014
		CN 103988150 A	13 Aug 2014
		CN 105122790 A	02 Dec 2015
		EP 1851750 A2	07 Nov 2007
		EP 1851750 A4	25 Aug 2010
		EP 2150893 A2	10 Feb 2010
		EP 2150893 A4	22 Aug 2012
		EP 2163987 A2	17 Mar 2010
		EP 2163987 A3	23 Jan 2013
		EP 2266016 A2	29 Dec 2010
		EP 2266016 A4	29 Oct 2014
		EP 2304527 A1	06 Apr 2011
		EP 2304527 A4	27 Mar 2013
		EP 2338114 A1	29 Jun 2011
		EP 2338114 A4	05 Oct 2011
		EP 2350774 A1	03 Aug 2011
		EP 2350774 A4	05 Nov 2014
		EP 2427857 A1	14 Mar 2012
		EP 2427857 A4	26 Dec 2012

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/IL2016/050202

Patent document cited search report	Publication date	Patent family member(s)	Publication Date
	EP 2435986	A2	04 Apr 2012
	EP 2435986	A4	22 Apr 2015
	EP 2689413	A2	29 Jan 2014
	EP 2689413	A4	22 Jul 2015
	EP 2724337	A1	30 Apr 2014
	EP 2724337	A4	17 Jun 2015
	EP 2845084	A2	11 Mar 2015
	EP 2893421	A1	15 Jul 2015
	EP 2918074	A1	16 Sep 2015
	EP 2941739	A2	11 Nov 2015
	JP 2008530661	A	07 Aug 2008
	JP 5631535	B2	26 Nov 2014
	JP 2011523112	A	04 Aug 2011
	JP 5697590	B2	08 Apr 2015
	JP 2012526328	A	25 Oct 2012
	JP 5698733	B2	08 Apr 2015
	JP 2012528405	A	12 Nov 2012
	JP 5782431	B2	24 Sep 2015
	JP 2013047978	A	07 Mar 2013
	JP 5788853	B2	07 Oct 2015
	JP 2012506097	A	08 Mar 2012
	JP 5805537	B2	04 Nov 2015
	JP 2012502364	A	26 Jan 2012
	JP 5806615	B2	10 Nov 2015
	JP 2010525494	A	22 Jul 2010
	JP 5905662	B2	20 Apr 2016
	JP 2011525283	A	15 Sep 2011
	JP 2014221636	A	27 Nov 2014
	JP 2014222506	A	27 Nov 2014
	JP 2015525381	A	03 Sep 2015
	JP 2016504652	A	12 Feb 2016

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/IL2016/050202

Patent document cited search report	Publication date	Patent family member(s)	Publication Date
	KR 20070116794	A	11 Dec 2007
	KR 100960577	B1	03 Jun 2010
	KR 20100017321	A	16 Feb 2010
	KR 101545360	B1	19 Aug 2015
	KR 20100136993	A	29 Dec 2010
	KR 101550478	B1	04 Sep 2015
	KR 20110058871	A	01 Jun 2011
	KR 101616591	B1	28 Apr 2016
	KR 20110022057	A	04 Mar 2011
	KR 20110079839	A	08 Jul 2011
	KR 20120034637	A	12 Apr 2012
	KR 20120034672	A	12 Apr 2012
	KR 20140033025	A	17 Mar 2014
	KR 20140068855	A	09 Jun 2014
	KR 20150105308	A	16 Sep 2015
	NO 20074200	A	07 Nov 2007
	NO 20093236	A	15 Jan 2010
	US 2006187196	A1	24 Aug 2006
	US 7598942	B2	06 Oct 2009
	US 2010131549	A1	27 May 2010
	US 8370383	B2	05 Feb 2013
	US 2008271053	A1	30 Oct 2008
	US 8407725	B2	26 Mar 2013
	US 2010060570	A1	11 Mar 2010
	US 8531396	B2	10 Sep 2013
	US 2010053304	A1	04 Mar 2010
	US 8537111	B2	17 Sep 2013
	US 2010060576	A1	11 Mar 2010
	US 8537112	B2	17 Sep 2013
	US 2011018803	A1	27 Jan 2011
	US 8665213	B2	04 Mar 2014

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/IL2016/050202

Patent document cited search report	Publication date	Patent family member(s)	Publication Date
	US 2011025598	A1	03 Feb 2011
	US 8669939	B2	11 Mar 2014
	US 2010281440	A1	04 Nov 2010
	US 8681098	B2	25 Mar 2014
	US 2010280988	A1	04 Nov 2010
	US 8723795	B2	13 May 2014
	US 2010090947	A1	15 Apr 2010
	US 8830168	B2	09 Sep 2014
	US 2010090946	A1	15 Apr 2010
	US 8866740	B2	21 Oct 2014
	US 2014145929	A1	29 May 2014
	US 8890813	B2	18 Nov 2014
	US 2012326963	A1	27 Dec 2012
	US 8896531	B2	25 Nov 2014
	US 2012268364	A1	25 Oct 2012
	US 8941588	B2	27 Jan 2015
	US 2013076522	A1	28 Mar 2013
	US 8941589	B2	27 Jan 2015
	US 2013076617	A1	28 Mar 2013
	US 8941590	B2	27 Jan 2015
	US 2010127968	A1	27 May 2010
	US 9052970	B2	09 Jun 2015
	US 2010128062	A1	27 May 2010
	US 9063801	B2	23 Jun 2015
	US 2009231278	A1	17 Sep 2009
	US 9075441	B2	07 Jul 2015
	US 2015054729	A1	26 Feb 2015
	US 9317128	B2	19 Apr 2016
	US 2009278915	A1	12 Nov 2009
	US 2010066676	A1	18 Mar 2010
	US 2011025603	A1	03 Feb 2011

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/IL2016/050202

Patent document cited search report	Publication date	Patent family member(s)	Publication Date
	US	2013076616 A1	28 Mar 2013
	US	2013312009 A1	21 Nov 2013
	US	2013347006 A1	26 Dec 2013
	US	2014035805 A1	06 Feb 2014
	US	2014145932 A1	29 May 2014
	US	2014152546 A1	05 Jun 2014
	US	2014195988 A1	10 Jul 2014
	US	2014225822 A1	14 Aug 2014
	US	2014240231 A1	28 Aug 2014
	US	2014298269 A1	02 Oct 2014
	US	2014325373 A1	30 Oct 2014
	US	2015077326 A1	19 Mar 2015
	US	2015100593 A1	09 Apr 2015
	US	2015138086 A1	21 May 2015
	US	2015205364 A1	23 Jul 2015
	US	2015309581 A1	29 Oct 2015
	US	2015316993 A1	05 Nov 2015
	US	2015331497 A1	19 Nov 2015
	US	2015371082 A1	24 Dec 2015
	US	2015371083 A1	24 Dec 2015
	US	2016018896 A1	21 Jan 2016
	US	2016070755 A1	10 Mar 2016
	US	2016116984 A1	28 Apr 2016
	WO	2006086508 A2	17 Aug 2006
	WO	2006086508 A9	14 Dec 2006
	WO	2006086508 A3	23 Apr 2009
	WO	2008134452 A2	06 Nov 2008
	WO	2008134452 A3	30 Dec 2009
	WO	2009124181 A2	08 Oct 2009
	WO	2009124181 A3	30 Dec 2009
	WO	2009155465 A1	23 Dec 2009

Information on patent family members

PCT/IL2016/050202

Form PCT/ISA/210 (patent family annex) (January 2015)