



(51) International Patent Classification:
G06K 9/00 (2006.01)

(21) International Application Number:
PCT/US2013/021709

(22) International Filing Date:
16 January 2013 (16.01.2013)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/587,554 17 January 2012 (17.01.2012) US
13/414,485 7 March 2012 (07.03.2012) US
61/724,091 8 November 2012 (08.11.2012) US
13/724,357 21 December 2012 (21.12.2012) US

(71) Applicant: **LEAP MOTION, INC.** [US/US]; 333 Bryant Street, Suite LL150, San Francisco, CA 94107 (US).

(72) Inventor; and

(71) Applicant : **HOLZ, David** [US/US]; 219 Brannan Street, Apt. 12J, San Francisco, CA 94107 (US).

(74) Agent: **LIN, Yi-chia**; Bingham McCutchen LLP, 2020 K Street NW, Washington, DC 20006 (US).

(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, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, 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, 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, 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, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: SYSTEMS AND METHODS FOR CAPTURING MOTION IN THREE-DIMENSIONAL SPACE

(57) Abstract: Methods and systems for capturing motion and/or determining the shapes and positions of one or more objects in 3D space utilize cross-sections thereof. In various embodiments, images of the cross-sections are captured using at least one camera.



WO 2013/109608 A2

5 **SYSTEMS AND METHODS FOR CAPTURING MOTION**
 IN THREE-DIMENSIONAL SPACE

CROSS-REFERENCE TO RELATED APPLICATION

10 [0001] This application claims priority to U.S. Patent Application Nos. 13/414,485 (filed on March 7, 2012) and 13/724,357 (filed on December 21, 2012), and also claims priority to and the benefit of U.S. Provisional Patent Application Nos. 61/724,091 (filed on November 8, 2012) and 61/587,554 (filed on January 17, 2012). The foregoing applications are incorporated herein by reference in their entireties.

15 **FIELD OF THE INVENTION**

 [0002] The present invention relates, in general, to image analysis, and in particular embodiments to identifying shapes and capturing motions of objects in three-dimensional space.

BACKGROUND

20 [0003] Motion capture has numerous applications. For example, in filmmaking, digital models generated using motion capture can be used as the basis for the motion of computer-generated characters or objects. In sports, motion capture can be used by coaches to study an athlete's movements and guide the athlete toward improved body mechanics. In video games or virtual reality applications, motion capture can be used to allow a person to interact with a virtual
25 environment in a natural way, e.g., by waving to a character, pointing at an object, or performing an action such as swinging a golf club or baseball bat.

 [0004] The term "motion capture" refers generally to processes that capture movement of a subject in three-dimensional (3D) space and translate that movement into, for example, a digital model or other representation. Motion capture is typically used with complex subjects that have
30 multiple separately articulating members whose spatial relationships change as the subject moves. For instance, if the subject is a walking person, not only does the whole body move across space, but the position of arms and legs relative to the person's core or trunk are constantly shifting. Motion capture systems are typically interested in modeling this articulation.

 [0005] Most existing motion capture systems rely on markers or sensors worn by the subject
35 while executing the motion and/or on the strategic placement of numerous cameras in the environment to capture images of the moving subject from different angles. Such systems tend to be expensive to construct. In addition, markers or sensors worn by the subject can be cumbersome

5 and interfere with the subject's natural movement. Further, systems involving large numbers of cameras tend not to operate in real time, due to the volume of data that needs to be analyzed and correlated. Such considerations of cost, complexity and convenience have limited the deployment and use of motion capture technology.

[0006] Consequently, there is a need for an economical approach that captures the motion of
10 objects in real time without attaching sensors or markers thereto.

SUMMARY

[0001] Embodiments of the present invention relate to methods and systems for capturing motion and/or determining position of an object using small amounts of information. For example,
15 an outline of an object's shape, or silhouette, as seen from a particular vantage point can be used to define tangent lines to the object from that vantage point in various planes, referred to herein as "slices" of the object. Using as few as two different vantage points, four (or more) tangent lines from the vantage points to the object can be obtained for a given slice. From the tangent lines, it is possible to determine the position of the slice and to approximate its cross-section. The locations
20 of points on an object's surface in a particular slice can be determined directly (e.g., using a time-of-flight camera), and the position and shape of a cross-section of the object in the slice can be approximated by fitting an ellipse or other simple closed curve to the points. Positions and cross-sections determined for different slices can be correlated to construct a 3D model of the object, including its position and shape. A succession of images can be analyzed using the same technique
25 to model motion of the object. Motion of a complex object that has multiple separately articulating members (e.g., a human hand) can be modeled using techniques described herein.

[0002] In various embodiments, object cross-sections are obtained from, for example, reflections from the object or shadows cast by the object. In various embodiments, the 3D reflections or shadows captured using a camera are first sliced into multiple two-dimensional (2D)
30 cross-sectional images. The cross-sectional position and shape (or "intersection region") of the 3D objects in each 2D slice may be determined based on the positions of one or more light sources used to illuminate the objects and the captured reflections or shadows. The 3D structure of the object may then be reconstructed by assembling a collection of the intersection regions obtained in the 2D slices. The objective, in general, is to obtain either a unique ellipse describing the cross-
35 section of the object, or a subset of the parameters defining the cross-section (in which case the remaining parameters may be estimated). If there are more light sources than are necessary to determine the shape of the cross-section, some optimized subset of them may be utilized for

5 maximum accuracy. The light sources may emit at different wavelengths so that their individual contributions are more easily identified, or they may be turned on in sequence rather than simultaneously, or they may have different brightnesses.

[0003] In some embodiments, the 2D intersection regions are identified based on “true” intersection points — i.e., points within the volume defined by the intersection of all light beams, which volume includes the object. These true intersection points may determined by the light sources and reflections or shadows — e.g., based on the number of reflection or shadow regions that they lie within or the locations of the geometric projection points calculated based on the positions of the light sources. In one embodiment, the light sources are arranged, for example, in a line or a plane such that the true intersection points are determined without identifying the actual locations thereof; this reduces the computational complexity, thereby increasing the processing speed. In some embodiments, the intersection region is split into a number of smaller intersection regions that can individually represent at least a portion of the reflections or shadows in the scene. Because determining each of the smaller intersection regions is computationally simpler than determining the entire intersection region, the processing time for obtaining the entire intersection region assembled from the individual smaller intersection regions is reduced (even if the smaller intersection regions are determined sequentially rather than in parallel). In various embodiments, the number of small split intersection regions that need to be identified is reduced by setting a criteria number U equal to the greatest number of intersection points in any intersection region; only regions or combinations of regions having a number of intersection points exceeding the criteria number U are further processed to identify the intersection regions therein.

[0004] In some embodiments, an image coordinate system using, for example, an imaging grid is incorporated into the system to easily define locations of the reflections or shadows. In one implementation, the camera includes multiple color filters placed on the light sensors to generate multiple images, each corresponding to a different color filter. Application of the 2D approaches described above to the color-specific images may then determine both the locations and colors of the objects.

[0005] In some embodiments, the 2D cross-section regions are identified based on a vantage point defined by the position of an image-capturing camera and shadow edge points generated by light sources. At the vantage point, two light rays are detected; these light rays are transmitted from a left-edge tangent point and a right-edge tangent point of the cross-section, and define a viewed portion of the cross-section within the field of view of the camera. Two equations based on the positions of the two edge tangent points can partially determine the characteristic parameters of

5 a closed curve (e.g., an ellipse) approximating the contour of the object's cross-section. Additionally, each shadow edge point created by emitting light from a light source onto the cross-section can provide two equations, one based on the detected position of the shadow edge point and the other based on the light ray emitted from the light source to the shadow edge point on the cross-section. Utilizing a suitable number (e.g., one or a plurality) of light sources can provide sufficient

10 information to determine the characteristic parameters of the fitting ellipse, thereby identifying the position and size of the cross-section. Accordingly, a 3D model of the object can be reconstructed by correlating the determined positions and sizes of the cross-sections in the 2D slices. A succession of images can then be analyzed using the same technique to model motion of the object.

[0006] Accordingly, in one aspect, the invention pertains to a method of identifying a position

15 and shape of an object (e.g., a human, a human body part, or a handheld object such as a pencil or a scalpel) in 3D space. In representative embodiments, the method includes capturing an image generated by casting an output from one or more sources (e.g., a light source or a sonic source) onto the object; analyzing the image to computationally slice the object into multiple 2D slices, where each slice corresponds to a cross-section of the object; identifying shapes and positions of multiple

20 cross-sections of the object based at least in part on the image and a location of the one or more sources; and reconstructing the position and shape of the object in 3D space based at least in part on the multiple identified cross-sectional shapes and positions of the object. The position and shape of the object in 3D space may be reconstructed based on correlations between the multiple 2D slices.

[0007] In various embodiments, the cross-sectional shape and position of the object is

25 identified by selecting a collection of intersection points generated by analyzing a location of the one or more sources and positions of points in the image (e.g., a shadow of the object) associated with the 2D slice. The intersection points may be selected based on the total number source(s) employed. Alternatively, the intersection points may be selected based on locations of projection points associated with the intersection points, where the projection points are projections from the

30 intersection points onto the 2D slice (e.g., where the projection is dictated by the position(s) of the source(s)). In some embodiments, the method further includes splitting the cross-section of the object into multiple regions and using each region to generate one or more portions of the shadow image of the 2D slice, and identifying the regions based on the shadow image of the 2D slice and the location of the one or more sources. A region may be established or recognized if the number

35 of the intersection points is equal to or greater than a predetermined criteria number. Additionally, the intersection points may be selected based on the location of the source(s) and the size of the image cross-section. The image may include reflections from the object and the intersection points

5 may be selected based on time-of-flight data using a time-of-flight camera. In one implementation, the selected collection of intersection points in a first 2D slice is reused in a second 2D slice. In addition, the image may be generated by casting light from multiple light sources, aligned in a line or in a plane, onto the object.

[0008] In one embodiment, the method includes defining a 3D model of the object and
10 reconstructing the position and shape of the object in 3D space based on the 3D model. In another embodiment, the method includes defining coordinates of the image. In one implementation, the image is separated into multiple primary images each including a color; various colors on the object are identified based on the primary images.

[0009] In various embodiments, the method includes manipulating one or more virtual objects
15 displayed on a device based on the identified position and shape of the object. The device may be a head-mounted device or a TV. In one embodiment, the identified position and shape of the object is used to manipulate the virtual object via wireless cell phone communication. In some embodiments, the method further includes authenticating a user based on the detected shape of the object and/or the detected motion of the object and subsequent matching thereof to data in a
20 database record corresponding to the user.

[0010] Some embodiments utilize a single camera to capture an image generated by casting an
output from at least one source onto the object; analyzing the image to computationally slice the
object into a plurality of 2D slices, each of which corresponds to a cross-section of the object, based
at least in part on multiple edge points in the image (where an edge point may be, e.g., an
25 illuminated edge point — i.e., a point on the edge of the object that is detectable by the camera —
or a shadow edge point at the boundary of a shadow region, as more fully described below); and
reconstructing the position and shape of at least a portion of the object in 3D space based at least in
part on a plurality of the identified cross-sectional positions and sizes. The source(s) may be one or
more light sources — e.g., one, two, three, or more than three light-emitting diodes (LEDs). A
30 plurality of light sources may be operated in a pulsed fashion, whereby a plurality of the edge
points are generated sequentially.

[0011] In various embodiments, the edge points define a viewed portion of the cross-section
within which the portion of the cross-section is within a field of view of an image-capturing device
(e.g., a camera). Light rays cast from the edge points to the image-capturing device may be tangent
35 to the cross-section, and at least one shadow edge point may be created by emitting light from the
source(s) onto the object. The shadow edge point(s) may be defined by a boundary between a
shadow region and an illuminated region on the cross-section of the object.

5 [0012] Methods in accordance with the invention may further comprise defining a 3D model of the object and reconstructing the position and shape of the object in 3D space based on the 3D model. The position and shape of the object in 3D space may be reconstructed based on correlations between the plurality of the 2D slices.

[0013] In another aspect, the invention relates to a system for identifying a position and shape
10 of an object in 3D space. In various embodiments, the system includes one or more cameras (e.g., a time-of-flight camera) oriented toward a field of view; one or more sources (e.g., a light source or a sonic source) to direct illumination onto the object in the field of view; and an image analyzer coupled to the camera and the source and configured to operate the camera to capture one or more images of the object and identify a position and shape of the object in 3D space based on the
15 captured image and a location of the source.

[0014] In one implementation, the one or more light sources include multiple light sources each aligned in a line or in a plane. Additionally, the system may include multiple filters placed on light sensors of the camera to generate multiple images, each of which corresponds to a color filter. In one embodiment, the image analyzer is further configured to (i) slice the object into multiple 2D
20 slices each corresponding to a cross-section of the object, (ii) identify a shape and position of the object based at least in part on an image captured by the camera and a location of the one or more light source, and (iii) reconstruct the position and shape of the object in 3D space based at least in part on the multiple identified cross-sectional shapes and positions of the object. In some embodiments, the image analyzer is further configured to define a 3D model of the object and
25 reconstruct the position and shape of the object in 3D space based on the 3D model.

[0015] In various embodiments, the system further includes a secondary device (e.g., a head-mounted device or a mobile device) operatively connected to the system. The secondary device may be an authentication server for authenticating a user based on a shape and/or a jitter of the user's hand detected by the image analyzer.

30 [0016] Certain embodiments comprise a camera oriented toward a field of view; at least one source to direct illumination onto the object in the field of view; and an image analyzer coupled to the camera and the source(s). The image analyzer is configured to capture an image generated by casting an output from at least one source onto the object analyze the image to computationally slice the object into a plurality of two-dimensional 2D slices, each of which corresponds to a cross-
35 section of the object, based at least in part on edge points in the image; and reconstruct the position and shape of at least a portion of the object in 3D space based at least in part on a plurality of the identified cross-sectional positions and sizes. The source(s) may be a plurality of light sources,

5 e.g., one, two, three or more LEDs. The system may include a driver for operating the sources in a pulsed fashion, whereby a plurality of the shadow edge points are generated sequentially. In some embodiments, the image analyzer is further configured to define a 3D model of the object and reconstruct the position and shape of the object in 3D space based on the 3D model.

[0017] Reference throughout this specification to “one example,” “an example,” “one
10 embodiment,” or “an embodiment” means that a particular feature, structure, or characteristic described in connection with the example is included in at least one example of the present technology. Thus, the occurrences of the phrases “in one example,” “in an example,” “one
embodiment,” or “an embodiment” in various places throughout this specification are not
necessarily all referring to the same example. Furthermore, the particular features, structures,
15 routines, steps, or characteristics may be combined in any suitable manner in one or more examples of the technology. The headings provided herein are for convenience only and are not intended to limit or interpret the scope or meaning of the claimed technology.

BRIEF DESCRIPTION OF THE DRAWINGS

20 [0018] In the drawings, like reference characters generally refer to the same parts throughout the different views. Also, the drawings are not necessarily to scale, with an emphasis instead generally being placed upon illustrating the principles of the invention. In the following description, various embodiments of the present invention are described with reference to the following drawings, in which:

25 [0019] FIG. 1 is a simplified illustration of a motion capture system according to an embodiment of the present invention;

[0020] FIG. 2 is a simplified block diagram of a computer system that can be used according to an embodiment of the present invention;

[0021] FIGs. 3A (top view) and 3B (side view) are conceptual illustrations of how slices are
30 defined in a field of view according to an embodiment of the present invention;

[0022] FIGs. 4A-4C are top views illustrating an analysis that can be performed on a given slice according to an embodiment of the present invention. FIG. 4A is a top view of a slice. FIG. 4B illustrates projecting edge points from an image plane to a vantage point to define tangent lines. FIG. 4C illustrates fitting an ellipse to tangent lines as defined in FIG. 4B;

35 [0023] FIG. 5 graphically illustrates an ellipse in the xy plane characterized by five parameters;

[0024] FIGs. 6A and 6B provide a flow diagram of a motion-capture process according to an embodiment of the present invention;

- 5 [0025] FIG. 7 graphically illustrates a family of ellipses that can be constructed from four tangent lines;
- [0026] FIG. 8 sets forth a general equation for an ellipse in the xy plane;
- [0027] FIG. 9 graphically illustrates how a centerline can be found for an intersection region with four tangent lines according to an embodiment of the present invention;
- 10 [0028] FIGs. 10A- 10N set forth equations that can be solved to fit an ellipse to four tangent 15 lines according to an embodiment of the present invention;
- [0029] FIGs. 11A-11C are top views illustrating instances of slices containing multiple disjoint cross-sections according to various embodiments of the present invention;
- [0030] FIG. 12 graphically illustrates a model of a hand that can be generated using a motion 15 capture system according to an embodiment of the present invention;
- [0031] FIG. 13 is a simplified system diagram for a motion-capture system with three cameras according to an embodiment of the present invention;
- [0032] FIG. 14 illustrates a cross-section of an object as seen from three vantage points in the system of FIG. 13;
- 20 [0033] FIG. 15 graphically illustrates a technique that can be used to find an ellipse from at least five tangents according to an embodiment of the present invention;
- [0034] FIGs. 16A, 16B, and 16C are simplified illustrations of a motion-capture system in accordance with an embodiment of the present invention;
- [0035] FIG. 17 schematically illustrates a system for capturing shadows of an object according 25 to an embodiment of the present invention;
- [0036] FIG. 18 schematically illustrates an ambiguity that can occur in the system of FIG. 17;
- [0037] FIG. 19 schematically illustrates another system for capturing shadows of an object according to another embodiment of the present invention;
- [0038] FIG. 20 graphically depicts a collection of the intersection regions defined by a virtual 30 rubber band stretched around multiple intersection regions in accordance with an embodiment of the invention;
- [0039] FIG. 21 schematically illustrates a simple intersection region constructed using two light sources in accordance with an embodiment of the invention;
- [0040] FIGs. 22A, 22B and 22C schematically depict determinations of true intersection points 35 in accordance with various embodiments of the invention;
- [0041] FIG. 23 schematically depicts an intersection region uniquely identified using a group of the intersection points;

5 [0042] FIG. 24 illustrates an image coordinate system incorporated to define the locations of the shadows in accordance with an embodiment of the invention;

[0043] FIG. 25A illustrates separate color images captured using color filters in accordance with an embodiment of the invention;

[0044] FIG. 25B depicts a reconstructed 3D image of the object;

10 [0045] FIGs. 26A, 26B, and 26C schematically illustrate a system for capturing an image of both the object and one or more shadows cast by the object from one or more light sources at known positions according to an embodiment of the present invention;

[0046] FIG. 27 schematically illustrates a camera-and-beamsplitter setup for a motion capture system according to another embodiment of the present invention;

15 [0047] FIG. 28 schematically illustrates a camera-and-pinhole setup for a motion capture system according to another embodiment of the present invention; and

[0048] FIGs. 29A, 29B, and 29C depict a motion capture system operatively connected to a head-mounted device, a mobile device, and an authentication server, respectively.

20

DETAILED DESCRIPTION

[0049] Embodiments of the present invention relate to methods and systems for capturing motion and/or determining position of an object using small amounts of information. For example, an outline of an object's shape, or silhouette, as seen from a particular vantage point can be used to define tangent lines to the object from that vantage point in various planes, referred to herein as
25 "slices." Using as few as two different vantage points, four (or more) tangent lines from the vantage points to the object can be obtained in a given slice. From these four (or more) tangent lines, it is possible to determine the position of the object in the slice and to approximate its cross-section in the slice, e.g., using one or more ellipses or other simple closed curves. As another example, locations of points on an object's surface in a particular slice can be determined directly (e.g., using
30 a time-of-flight camera), and the position and shape of a cross-section of the object in the slice can be approximated by fitting an ellipse or other simple closed curve to the points. Positions and cross-sections determined for different slices can be correlated to construct a 3D model of the object, including its position and shape. A succession of images can be analyzed using the same technique to model motion of the object. Motion of a complex object that has multiple separately articulating
35 members (e.g., a human hand) can be modeled using techniques described herein.

[0050] In some embodiments, the silhouettes of an object are extracted from one or more images of the object that reveal information about the object as seen from different vantage points.

5 While silhouettes can be obtained using a number of different techniques, in some embodiments, the silhouettes are obtained by using cameras to capture images of the object and analyzing the images to detect object edges.

[0051] FIG. 1 is a simplified illustration of a motion capture system 100 according to an embodiment of the present invention. System 100 includes two cameras 102, 104 arranged such
10 that their fields of view (indicated by broken lines) overlap in region 110. Cameras 102 and 104 are coupled to provide image data to a computer 106. Computer 106 analyzes the image data to determine the 3D position and motion of an object, e.g., a hand 108., that moves in the field of view of cameras 102, 104.

[0052] Cameras 102, 104 can be any type of camera, including visible-light cameras, infrared
15 (IR) cameras, ultraviolet cameras or any other devices (or combination of devices) that are capable of capturing an image of an object and representing that image in the form of digital data. Cameras 102, 104 are preferably capable of capturing video images (i.e., successive image frames at a constant rate of at least 15 frames per second), although no particular frame rate is required. The particular capabilities of cameras 102, 104 are not critical to the invention, and the cameras can
20 vary as to frame rate, image resolution (e.g., pixels per image), color or intensity resolution (e.g., number of bits of intensity data per pixel), focal length of lenses, depth of field, etc. In general, for a particular application, any cameras capable of focusing on objects within a spatial volume of interest can be used. For instance, to capture motion of the hand of an otherwise stationary person, the volume of interest might be a meter on a side. To capture motion of a running person, the
25 volume of interest might be tens of meters in order to observe several strides (or the person might run on a treadmill, in which case the volume of interest can be considerably smaller).

[0053] The cameras can be oriented in any convenient manner. In the embodiment shown, respective optical axes 112, 114 of cameras 102 and 104 are parallel, but this is not required. As described below, each camera is used to define a “vantage point” from which the object is seen, and
30 it is required only that a location and view direction associated with each vantage point be known, so that the locus of points in space that project onto a particular position in the camera’s image plane can be determined. In some embodiments, motion capture is reliable only for objects in area 110 (where the fields of view of cameras 102, 104 overlap), and cameras 102, 104 may be arranged to provide overlapping fields of view throughout the area where motion of interest is expected to
35 occur.

[0054] In FIG. 1 and other examples described herein, object 108 is depicted as a hand. The hand is used only for purposes of illustration, and it is to be understood that any other object can be

5 the subject of motion capture analysis as described herein. Computer 106 can be any device that is capable of processing image data using techniques described herein. FIG. 2 is a simplified block diagram of computer system 200 implementing computer 106 according to an embodiment of the present invention. Computer system 200 includes a processor 202, a memory 204, a camera interface 206, a display 208, speakers 209, a keyboard 210, and a mouse 211.

10 [0055] Processor 202 can be of generally conventional design and can include, e.g., one or more programmable microprocessors capable of executing sequences of instructions. Memory 204 can include volatile (e.g., DRAM) and nonvolatile (e.g., flash memory) storage in any combination. Other storage media (e.g., magnetic disk, optical disk) can also be provided. Memory 204 can be used to store instructions to be executed by processor 202 as well as input and/or output data
15 associated with execution of the instructions.

[0056] Camera interface 206 can include hardware and/or software that enables communication between computer system 200 and cameras such as cameras 102, 104 of FIG. 1. Thus, for example, camera interface 206 can include one or more data ports 216, 218 to which cameras can be connected, as well as hardware and/or software signal processors to modify data signals received
20 from the cameras (e.g., to reduce noise or reformat data) prior to providing the signals as inputs to a conventional motion-capture (“mocap”) program 214 executing on processor 202. In some embodiments, camera interface 206 can also transmit signals to the cameras, e.g., to activate or deactivate the cameras, to control camera settings (frame rate, image quality, sensitivity, etc.), or the like. Such signals can be transmitted, e.g., in response to control signals from processor 202,
25 which may in turn be generated in response to user input or other detected events.

[0057] In some embodiments, memory 204 can store mocap program 214, which includes instructions for performing motion capture analysis on images supplied from cameras connected to camera interface 206. In one embodiment, mocap program 214 includes various modules, such as an image analysis module 222, a slice analysis module 224, and a global analysis module 226.

30 Image analysis module 222 can analyze images, e.g., images captured via camera interface 206, to detect edges or other features of an object. Slice analysis module 224 can analyze image data from a slice of an image as described below, to generate an approximate cross-section of the object in a particular plane. Global analysis module 226 can correlate cross-sections across different slices and refine the analysis. Examples of operations that can be implemented in code modules of mocap
35 program 214 are described below.

[0058] Memory 204 can also include other information used by mocap program 214; for example, memory 204 can store image data 228 and an object library 230 that can include

5 canonical models of various objects of interest. As described below, an object being modeled can be identified by matching its shape to a model in object library 230.

[0059] Display 208, speakers 209, keyboard 210, and mouse 211 can be used to facilitate user interaction with computer system 200. These components can be of generally conventional design or modified as desired to provide any type of user interaction. In some embodiments, results of motion capture using camera interface 206 and mocap program 214 can be interpreted as user input. For example, a user can perform hand gestures that are analyzed using mocap program 214, and the results of this analysis can be interpreted as an instruction to some other program executing on processor 200 (e.g., a web browser, word processor or the like). Thus, by way of illustration, a user might be able to use upward or downward swiping gestures to “scroll” a webpage currently displayed on display 208, to use rotating gestures to increase or decrease the volume of audio output from speakers 209, and so on.

[0060] It will be appreciated that computer system 200 is illustrative and that variations and modifications are possible. Computers can be implemented in a variety of form factors, including server systems, desktop systems, laptop systems, tablets, smart phones or personal digital assistants, and so on. A particular implementation may include other functionality not described herein, e.g., wired and/or wireless network interfaces, media playing and/or recording capability, etc. In some embodiments, one or more cameras may be built into the computer rather than being supplied as separate components.

[0061] While computer system 200 is described herein with reference to particular blocks, it is to be understood that the blocks are defined for convenience of description and are not intended to imply a particular physical arrangement of component parts. Further, the blocks need not correspond to physically distinct components. To the extent that physically distinct components are used, connections between components (e.g., for data communication) can be wired and/or wireless as desired.

[0062] An example of a technique for motion capture using the system of FIGs. 1 and 2 will now be described. In this embodiment, cameras 102, 104 are operated to collect a sequence of images of an object 108. The images are time correlated such that an image from camera 102 can be paired with an image from camera 104 that was captured at the same time (within a few milliseconds). These images are then analyzed, e.g., using mocap program 214, to determine the object’s position and shape in 3D space. In some embodiments, the analysis considers a stack of 2D cross-sections through the 3D spatial field of view of the cameras. These cross-sections are referred to herein as “slices.”

5 [0063] FIGs. 3A and 3B are conceptual illustrations of how slices are defined in a field of view according to an embodiment of the present invention. FIG. 3A shows, in top view, cameras 102 and 104 of FIG. 1. Camera 102 defines a vantage point 302, and camera 104 defines a vantage point 304. Line 306 joins vantage points 302 and 304. FIG. 3B shows a side view of cameras 102 and 104; in this view, camera 104 happens to be directly behind camera 102 and thus occluded; line
10 306 is perpendicular to the plane of the drawing. (It should be noted that the designation of these views as “top” and “side” is arbitrary; regardless of how the cameras are actually oriented in a particular setup, the “top” view can be understood as a view looking along a direction normal to the plane of the cameras, while the “side” view is a view in the plane of the cameras.)

[0064] An infinite number of planes can be drawn through line 306. A “slice” can be any one of
15 those planes for which at least part of the plane is in the field of view of cameras 102 and 104. Several slices 308 are shown in FIG. 3B. (Slices 308 are seen edge-on; it is to be understood that they are 2D planes and not 1-D lines.) For purposes of motion capture analysis, slices can be selected at regular intervals in the field of view. For example, if the received images include a fixed number of rows of pixels (e.g., 1080 rows), each row can be a slice, or a subset of the rows can be
20 used for faster processing. Where a subset of the rows is used, image data from adjacent rows can be averaged together, e.g., in groups of 2-3.

[0065] FIGs. 4A-4C illustrate an analysis that can be performed on a given slice. FIG. 4A is a top view of a slice as defined above, corresponding to an arbitrary cross-section 402 of an object. Regardless of the particular shape of cross-section 402, the object as seen from a first vantage point
25 404 has a “left illuminated edge” point 406 and a “right illuminated edge” point 408. As seen from a second vantage point 410, the same object has a “left illuminated edge” point 412 and a “right illuminated edge” point 414. These are in general different points on the boundary of object 402. A tangent line can be defined that connects each illuminated edge point and the associated vantage point. For example, FIG. 4A also shows that tangent line 416 can be defined through vantage point
30 404 and left illuminated edge point 406; tangent line 418 through vantage point 404 and right illuminated edge point 408; tangent line 420 through vantage point 410 and left illuminated edge point 412; and tangent line 422 through vantage point 410 and right illuminated edge point 414.

[0066] It should be noted that all points along any one of tangent lines 416, 418, 420, 422 will project to the same point on an image plane. Therefore, for an image of the object from a given
35 vantage point, a left illuminated edge point and a right illuminated edge point can be identified in the image plane and projected back to the vantage point, as shown in FIG. 4B, which is another top view of a slice, showing the image plane for each vantage point. Image 440 is obtained from

vantage point 442 and shows left illuminated edge point 446 and right illuminated edge point 448. Image 450 is obtained from vantage point 452 and shows left illuminated edge point 456 and right illuminated edge point 458. Tangent lines 462, 464, 466, 468 can be defined as shown. Given the tangent lines of FIG. 4B, the location in the slice of an elliptical cross-section can be determined, as illustrated in FIG. 4C, where ellipse 470 has been fit to tangent lines 462, 464, 466, 468 of FIG. 4B.

[0067] In general, as shown in FIG. 5, an ellipse in the xy plane can be characterized by five parameters: the x and y coordinates of the center (x_C , y_C), the semimajor axis (a), the semiminor axis (b), and a rotation angle (θ) (e.g., the angle of the semimajor axis relative to the x axis). With only four tangents, as is the case in FIG. 4C, the ellipse is underdetermined. However, an efficient process for estimating the ellipse in spite of this has been developed. In various embodiments as described below, this involves making an initial working assumption (or “guess”) as to one of the parameters and revisiting the assumption as additional information is gathered during the analysis. This additional information can include, for example, physical constraints based on properties of the cameras and/or the object.

[0068] In some embodiments, more than four tangents to an object may be available for some or all of the slices, e.g., because more than two vantage points are available. An elliptical cross-section can still be determined, and the process in some instances is somewhat simplified as there is no need to assume a parameter value. In some instances, the additional tangents may create additional complexity. Examples of processes for analysis using more than four tangents are described below and in the '554 application noted above.

[0069] In some embodiments, fewer than four tangents to an object may be available for some or all of the slices, e.g., because an edge of the object is out of range of the field of view of one camera or because an edge was not detected. A slice with three tangents can be analyzed. For example, using two parameters from an ellipse fit to an adjacent slice (e.g., a slice that had at least four tangents), the system of equations for the ellipse and three tangents is sufficiently determined that it can be solved. As another option, a circle can be fit to the three tangents; defining a circle in a plane requires only three parameters (the center coordinates and the radius), so three tangents suffice to fit a circle. Slices with fewer than three tangents can be discarded or combined with adjacent slices.

[0070] In some embodiments, each of a number of slices is analyzed separately to determine the size and location of an elliptical cross-section of the object in that slice. This provides an initial 3D model (specifically, a stack of elliptical cross-sections), which can be refined by correlating the cross-sections across different slices. For example, it is expected that an object's surface will have

5 continuity, and discontinuous ellipses can accordingly be discounted. Further refinement can be obtained by correlating the 3D model with itself across time, e.g., based on expectations related to continuity in motion and deformation.

[0071] A further understanding of the analysis process can be had by reference to FIGs. 6A-6B, which provide a flow diagram of a motion-capture process 600 according to an embodiment of the
10 present invention. Process 600 can be implemented, e.g., in mocap program 214 of FIG. 2.

[0072] At block 602, a set of images — e.g., one image from each camera 102, 104 of FIG. 1 — is obtained. In some embodiments, the images in a set are all taken at the same time (or within a few milliseconds), although a precise timing is not required. The techniques described herein for constructing an object model assume that the object is in the same place in all images in a set,
15 which will be the case if images are taken at the same time. To the extent that the images in a set are taken at different times, motion of the object may degrade the quality of the result, but useful results can be obtained as long as the time between images in a set is small enough that the object does not move far, with the exact limits depending on the particular degree of precision desired.

[0073] At block 604, each slice is analyzed. FIG. 6B illustrates a per-slice analysis that can be
20 performed at block 604. Referring to FIG. 6B, at block 606, illuminated edge points of the object in a given slice are identified in each image in the set. For example, edges of an object in an image can be detected using conventional techniques, such as contrast between adjacent pixels or groups of pixels. In some embodiments, if no illuminated edge points are detected for a particular slice (or if only one illuminated edge point is detected), no further analysis is performed on that slice. In
25 some embodiments, edge detection can be performed for the image as a whole rather than on a per-slice basis.

[0074] At block 608, assuming enough illuminated edge points were identified, a tangent line from each illuminated edge point to the corresponding vantage point is defined, e.g., as shown in FIG. 4C and described above. At block 610 an initial assumption as to the value of one of the
30 parameters of an ellipse is made, to reduce the number of free parameters from five to four. In some embodiments, the initial assumption can be, e.g., the semimajor axis (or width) of the ellipse. Alternatively, an assumption can be made as to eccentricity (ratio of semimajor axis to semiminor axis), and that assumption also reduces the number of free parameters from five to four. The assumed value can be based on prior information about the object. For example, if previous
35 sequential images of the object have already been analyzed, it can be assumed that the dimensions of the object do not significantly change from image to image. As another example, if it is assumed that the object being modeled is a particular type of object (e.g., a hand), a parameter value can be

5 assumed based on typical dimensions for objects of that type (e.g., an average cross-sectional dimension of a palm or finger). An arbitrary assumption can also be used, and any assumption can be refined through iterative analysis as described below.

[0075] At block 612, the tangent lines and the assumed parameter value are used to compute the other four parameters of an ellipse in the plane. For example, as shown in FIG. 7, four tangent
 10 lines 701, 702, 703, 704 define a family of inscribed ellipses 706 including ellipses 706a, 706b, and 706c, where each inscribed ellipse 706 is tangent to all four of lines 701-704. Ellipse 706a and 706b represent the “extreme” cases (i.e., the most eccentric ellipses that are tangent to all four of lines 701-704. Intermediate between these extremes are an infinite number of other possible ellipses, of which one example, ellipse 706c, is shown (dashed line).

15 [0076] The solution process selects one (or in some instances more than one) of the possible inscribed ellipses 706. In one embodiment, this can be done with reference to the general equation for an ellipse shown in FIG. 8. The notation follows that shown in FIG. 5, with (x, y) being the coordinates of a point on the ellipse, (x_C, y_C) the center, a and b the axes, and θ the rotation angle. The coefficients C_1 , C_2 and C_3 are defined in terms of these parameters, as shown in FIG. 8.

20 [0077] The number of free parameters can be reduced based on the observation that the centers (x_C, y_C) of all the ellipses in family 706 line on a line segment 710 (also referred to herein as the “centerline”) between the center of ellipse 706a (shown as point 712a) and the center of ellipse 706b (shown as point 712b). FIG. 9 illustrates how a centerline can be found for an intersection region. Region 902 is a “closed” intersection region; that is, it is bounded by tangents 904, 906,
 25 908, 910. The centerline can be found by identifying diagonal line segments 912, 914 that connect the opposite corners of region 902, identifying the midpoints 916, 918 of these line segments, and identifying the line segment 920 joining the midpoints as the centerline.

[0078] Region 930 is an “open” intersection region; that is, it is only partially bounded by tangents 904, 906, 908, 910. In this case, only one diagonal, line segment 932, can be defined. To
 30 define a centerline for region 930, centerline 920 from closed intersection region 902 can be extended into region 930 as shown. The portion of extended centerline 920 that is beyond line segment 932 is centerline 940 for region 930. In general, for any given set of tangent lines, both region 902 and region 930 can be considered during the solution process. (Often, one of these regions is outside the field of view of the cameras and can be discarded at a later stage.) Defining
 35 the centerline reduces the number of free parameters from five to four because y_C can be expressed as a (linear) function of x_C (or vice versa), based solely on the four tangent lines. However, for every point (x_C, y_C) on the centerline, a set of parameters $\{\theta, a, b\}$ can be found for an inscribed

5 ellipse. To reduce this to a set of discrete solutions, an assumed parameter value can be used. For example, it can be assumed that the semimajor axis a has a fixed value a_0 . Then, only solutions $\{\theta, a, b\}$ that satisfy $a = a_0$ are accepted.

[0079] In one embodiment, the ellipse equation of FIG. 8 is solved for θ , subject to the constraints that: (1) (x_C, y_C) must lie on the centerline determined from the four tangents (i.e., either centerline 920 or centerline 940 of FIG. 9); and (2) a is fixed at the assumed value a_0 . The ellipse equation can either be solved for θ analytically or solved using an iterative numerical solver (e.g., a Newtonian solver as is known in the art). An analytic solution can be obtained by writing an equation for the distances to the four tangent lines given a y_C position, then solving for the value of y_C that corresponds to the desired radius parameter $a = a_0$. One analytic solution is illustrated in the equations of FIGs. 10A-10D. Shown in FIG. 10A are equations for four tangent lines in the xy plane (the slice). Coefficients A_i, B_i and D_i (for $i=1$ to 4) can be determined from the tangent lines identified in an image slice as described above. FIG. 10B illustrates the definition of four column vectors $\mathbf{r}_{12}, \mathbf{r}_{23}, \mathbf{r}_{14}$ and $\mathbf{r}_{24}$ from the coefficients of FIG. 10A. The “ $\backslash$ ” operator here denotes matrix left division, which is defined for a square matrix $\mathbf{M}$ and a column vector $\mathbf{v}$ such that $\mathbf{M} \backslash \mathbf{v} = \mathbf{r}$, where $\mathbf{r}$ is the column vector that satisfies $\mathbf{M}\mathbf{r} = \mathbf{v}$. FIG. 10C illustrates the definition of $\mathbf{G}$ and $\mathbf{H}$, which are four-component vectors from the vectors of tangent coefficients $\mathbf{A}, \mathbf{B}$ and $\mathbf{D}$ and scalar quantities p and q , which are defined using the column vectors $\mathbf{r}_{12}, \mathbf{r}_{23}, \mathbf{r}_{14}$ and $\mathbf{r}_{24}$ from FIG. 10B. FIG. D illustrates the definition of six scalar quantities $v_{A2}, v_{AB}, v_{B2}, w_{A2}, w_{AB}$, and w_{B2} in terms of the components of vectors $\mathbf{G}$ and $\mathbf{H}$ of FIG. 10C.

25 [0080] Using the parameters defined in FIGs. 10A-10D, solving for θ is accomplished by solving the eighth-degree polynomial equation shown in FIG. 10E for t , where the coefficients Q_i (for $i = 0$ to 8) are defined as shown in FIGs. 10F-10N. The parameters $A_1, B_1, G_1, H_1, v_{A2}, v_{AB}, v_{B2}, w_{A2}, w_{AB}$, and w_{B2} used in FIGs. 10F-10N are defined as shown in FIGs. 10A-10D. The parameter n is the assumed semimajor axis (in other words, a_0). Once the real roots t are known, the possible values of θ are defined as $\theta = \text{atan}(t)$.

[0081] As it happens, the equation of FIGs. 10E-10N has at most three real roots; thus, for any four tangent lines, there are at most three possible ellipses that are tangent to all four lines and satisfy the $a = a_0$ constraint. (In some instances, there may be fewer than three real roots.) For each real root θ , the corresponding values of (x_C, y_C) and b can be readily determined. Depending on the particular inputs, zero or more solutions will be obtained; for example, in some instances, three solutions can be obtained for a typical configuration of tangents. Each solution is completely characterized by the parameters $\{\theta, a = a_0, b, (x_C, y_C)\}$.

5 [0082] Referring again to FIG. 6B, at block 614, the solutions are filtered by applying various constraints based on known (or inferred) physical properties of the system. For example, some solutions would place the object outside the field of view of the cameras, and such solutions can readily be rejected. As another example, in some embodiments, the type of object being modeled is known (e.g., it can be known that the object is or is expected to be a human hand). Techniques for
10 determining object type are described below; for now, it is noted that where the object type is known, properties of that object can be used to rule out solutions where the geometry is inconsistent with objects of that type. For example, human hands have a certain range of sizes and expected eccentricities in various cross-sections, and such ranges can be used to filter the solutions in a particular slice. These constraints can be represented in any suitable format, e.g., a physical
15 model (as described below), an ordered list of parameters based on such a model, etc.

[0083] In some embodiments, cross-slice correlations can also be used to filter (or further filter) the solutions obtained at block 612. For example, if the object is known to be a hand, constraints on the spatial relationship between various parts of the hand (e.g., fingers have a limited range of motion relative to each other and/or to the palm of the hand) as represented in a physical model or
20 explicit set of constraint parameters can be used to constrain one slice based on results from other slices. For purposes of cross-slice correlations, it should be noted that, as a result of the way slices are defined, the various slices may be tilted relative to each other, e.g., as shown in FIG. 3B. Accordingly, each planar cross-section can be further characterized by an additional angle θ , which can be defined relative to a reference direction 310 as shown in FIG. 3B.

25 [0084] At block 616, it is determined whether a satisfactory solution has been found. Various criteria can be used to assess whether a solution is satisfactory. For instance, if a unique solution is found (after filtering), that solution can be accepted, in which case process 600 proceeds to block 620 (described below). If multiple solutions remain or if all solutions were rejected in the filtering at block 614, it may be desirable to retry the analysis. If so, process 600 can return to block 610,
30 allowing a change in the assumption used in computing the parameters of the ellipse.

[0085] Retrying can be triggered under various conditions. For example, in some instances, 30 the initial parameter assumption (e.g., $a = a_0$) may produce no solutions or only nonphysical solutions (e.g., object outside the cameras' field of view). In this case, the analysis can be retried with a different assumption. In one embodiment, a small constant (which can be positive or
35 negative) is added to the initial assumed parameter value (e.g., a_0) and the new value is used to generate a new set of solutions. This can be repeated until an acceptable solution is found (or until the parameter value reaches a limit). An alternative approach is to keep the same assumption but to

5 relax the constraint that the ellipse be tangent to all four lines, e.g., by allowing the ellipse to be nearly but not exactly tangent to one or more of the lines. (In some embodiments, this relaxed constraint can also be used in the initial pass through the analysis.)

[0086] It should be noted that in some embodiments, multiple elliptical cross-sections may be found in some or all of the slices. For example, in some planes, a complex object (e.g., a hand) may
10 have a cross-section with multiple disjoint elements (e.g., in a plane that intersects the fingers). Ellipse-based reconstruction techniques as described herein can account for such complexity; examples are described below. Thus, it is generally not required that a single ellipse be found in a slice, and in some instances, solutions entailing multiple ellipses may be favored.

[0087] For a given slice, the analysis of FIG. 6B yields zero or more elliptical cross-sections. In
15 some instances, even after filtering at block 616, there may still be two or more possible solutions. These ambiguities can be addressed in further processing as described below.

[0088] Referring again to FIG. 6A, the per-slice analysis of block 604 can be performed for any number of slices, and different slices can be analyzed in parallel or sequentially, depending on available processing resources. The result is a 3D model of the object, where the model is
20 constructed by, in effect, stacking the slices. At block 620, cross-slice correlations are used to refine the model. For example, as noted above, in some instances, multiple solutions may have been found for a particular slice. It is likely that the “correct” solution (i.e., the ellipse that best corresponds to the actual position of the object) will correlate well with solutions in other slices, while any “spurious” solutions (i.e., ellipses that do not correspond to the actual position of the
25 object) will not. Uncorrelated ellipses can be discarded. In some embodiments where slices are analyzed sequentially, block 620 can be performed iteratively as each slice is analyzed.

[0089] At block 622, the 3D model can be further refined, e.g., based on an identification of the type of object being modeled. In some embodiments, a library of object types can be provided (e.g., as object library 230 of FIG. 2). For each object type, the library can provide characteristic
30 parameters for the object in a range of possible poses (e.g., in the case of a hand, the poses can include different finger positions, different orientations relative to the cameras, etc.). Based on these characteristic parameters, a reconstructed 3D model can be compared to various object types in the library. If a match is found, the matching object type is assigned to the model.

[0090] Once an object type is determined, the 3D model can be refined using constraints based
35 on characteristics of the object type. For instance, a human hand would characteristically have five fingers (not six), and the fingers would be constrained in their positions and angles relative to each other and to a palm portion of the hand. Any ellipses in the model that are inconsistent with these

5 constraints can be discarded. In some embodiments, block 622 can include recomputing all or portions of the per-slice analysis (block 604) and/or cross-slice correlation analysis (block 620) subject to the type-based constraints. In some instances, applying type-based constraints may cause deterioration in accuracy of reconstruction if the object is misidentified. (Whether this is a concern depends on implementation, and type-based constraints can be omitted if desired.)

10 [0091] In some embodiments, object library 230 can be dynamically and/or iteratively updated. For example, based on characteristic parameters, an object being modeled can be identified as a hand. As the motion of the hand is modeled across time, information from the model can be used to revise the characteristic parameters and/or define additional characteristic parameters, e.g., additional poses that a hand may present.

15 [0092] In some embodiments, refinement at block 622 can also include correlating results of analyzing images across time. It is contemplated that a series of images can be obtained as the object moves and/or articulates. Since the images are expected to include the same object, information about the object determined from one set of images at one time can be used to constrain the model of the object at a later time. (Temporal refinement can also be performed
20 “backward” in time, with information from later images being used to refine analysis of images at earlier times.)

[0093] At block 624, a next set of images can be obtained, and process 600 can return to block 604 to analyze slices of the next set of images. In some embodiments, analysis of the next set of images can be informed by results of analyzing previous sets. For example, if an object type was
25 determined, type-based constraints can be applied in the initial per-slice analysis, on the assumption that successive images are of the same object. In addition, images can be correlated across time, and these correlations can be used to further refine the model, e.g., by rejecting discontinuous jumps in the object’s position or ellipses that appear at one time point but completely disappear at the next.

30 [0094] It will be appreciated that the motion capture process described herein is illustrative and that variations and modifications are possible. Steps described as sequential may be executed in parallel, order of steps may be varied, and steps may be modified, combined, added or omitted. Different mathematical formulations and/or solution procedures can be substituted for those shown herein. Various phases of the analysis can be iterated, as noted above, and the degree to which
35 iterative improvement is used may be chosen based on a particular application of the technology. For example, if motion capture is being used to provide real-time interaction (e.g., to control a computer system), the data capture and analysis should be performed fast enough that the system

5 response feels like real time to the user. Inaccuracies in the model can be tolerated as long as they do not adversely affect the interpretation or response to a user's motion. In other applications, e.g., where the motion capture data is to be used for rendering in the context of digital movie-making, an analysis with more iterations that produces a more refined (and accurate) model may be preferred. As noted above, an object being modeled can be a "complex" object and consequently may present
10 multiple discrete ellipses in some cross-sections. For example, a hand has fingers, and a cross-section through the fingers may include as many as five discrete elements. The analysis techniques described above can be used to model complex objects.

[0095] By way of example, FIGs. 11A-11C illustrate some cases of interest. In FIG. 11A, cross-sections 1102, 1104 would appear as distinct objects in images from both of vantage points
15 1106, 1108. In some embodiments, it is possible to distinguish object from background; for example, in an infrared image, a heat-producing object (e.g., living organisms) may appear bright against a dark background. Where object can be distinguished from background, tangent lines 1110 and 1111 can be identified as a pair of tangents associated with opposite edges of one apparent object while tangent lines 1112 and 1113 can be identified as a pair of tangents associated with
20 opposite edges of another apparent object. Similarly, tangent lines 1114 and 1115, and tangent lines 1116 and 1117 can be paired. If it is known that vantage points 1106 and 1108 are on the same side of the object to be modeled, it is possible to infer that tangent pairs 1110, 1111 and 1116, 1117 should be associated with the same apparent object, and similarly for tangent pairs 1112, 1113 and 1114, 1115. This reduces the problem to two instances of the ellipse-fitting process described
25 above. If less information is available, an optimum solution can be determined by iteratively trying different possible assignments of the tangents in the slice in question, rejecting non-physical solutions, and cross-correlating results from other slices to determine the most likely set of ellipses.

[0096] In FIG. 11B, ellipse 1120 partially occludes ellipse 1122 from both vantage points. In some embodiments, it may or may not be possible to detect the "occlusion" edges 1124, 1126. If
30 edges 1124 and 1126 are not detected, the image appears as a single object and is reconstructed as a single elliptical cross-section. In this instance, information from other slices or temporal correlation across images may reveal the error. If occlusion edges 1124 and/or 1126 are visible, it may be apparent that there are multiple objects (or that the object has a complex shape) but it may not be apparent which object or object portion is in front. In this case, it is possible to compute multiple
35 alternative solutions, and the optimum solution may be ambiguous. Spatial correlations across slices, temporal correlations across image sets, and/or physical constraints based on object type can be used to resolve the ambiguity.

5 [0097] In FIG. 11C, ellipse 1140 fully occludes ellipse 1142. In this case, the analysis described above would not show ellipse 1142 in this particular slice. However, spatial correlations across slices, temporal correlations across image sets, and/or physical constraints based on object type can be used to infer the presence of ellipse 1142, and its position can be further constrained by the fact that it is apparently occluded. In some embodiments, multiple discrete cross-sections (e.g., in any
10 of FIGs. 11A-11C) can also be resolved using successive image sets across time. For example, the four-tangent slices for successive images can be aligned and used to define a slice with 5-8 tangents. This slice can be analyzed using techniques described below.

[0098] In one embodiment of the present invention, a motion capture system can be used to detect the 3D position and movement of a human hand. In this embodiment, two cameras are
15 arranged as shown in FIG. 1, with a spacing of about 1.5 cm between them. Each camera is an infrared camera with an image rate of about 60 frames per second and a resolution of 640x480 pixels per frame. An infrared light source (e.g., an IR light-emitting diode) that approximates a point light source is placed between the cameras to create a strong contrast between the object of interest (in this case, a hand) and background. The falloff of light with distance creates a strong
20 contrast if the object is a few inches away from the light source while the background is several feet away.

[0099] The image is analyzed using contrast between adjacent pixels to detect edges of the object. Bright pixels (detected illumination above a threshold) are assumed to be part of the object while dark pixels (detected illumination below a threshold) are assumed to be part of the
25 background. Edge detection may take approximately 2 ms with conventional processing capability. The edges and the known camera positions are used to define tangent lines in each of 480 slices (one slice per row of pixels), and ellipses are determined from the tangents using the analytical technique described above with reference to FIGs. 6A and 6B. In a typical case of modeling a hand, roughly 800-1200 ellipses are generated from a single pair of image frames (the number depends on
30 the orientation and shape of the hand) within, in various embodiments, about 6 ms. The error in modeling finger position in one embodiment is less than 0.1 mm.

[00100] FIG. 12 illustrates a model 1200 of a hand that can be generated using the system just described. As can be seen, the model does not have the exact shape of a hand, but a palm 1202, thumb 1204 and four fingers 1206 can be clearly recognized. Such models can be useful as the
35 basis for constructing more realistic models. For example, a skeleton model for a hand can be defined, and the positions of various joints in the skeleton model can be determined by reference to model 1200. Using the skeleton model, a more realistic image of a hand can be rendered.

5 Alternatively, a more realistic model may not be needed. For example, model 1200 accurately indicates the position of thumb 1204 and fingers 1206, and a sequence of models 1200 captured across time will indicate movement of these digits. Thus, gestures can be recognized directly from model 1200. The point is that ellipses identified and tracked as described above can be used to drive visual representations of the object tracked by application to a physical model of the object.

10 The model may be selected based on a desired degree of realism, the response time desired (or the latency that can be tolerated), and available computational resources.

[00101] It will be appreciated that this example system is illustrative and that variations and modifications are possible. Different types and arrangements of cameras can be used, and appropriate image analysis techniques can be used to distinguish object from background and thereby determine a silhouette (or a set of edge locations for the object) that can in turn be used to define tangent lines to the object in various 2D slices as described above. Given four tangent lines to an object, where the tangents are associated with at least two vantage points, an elliptical cross-section can be determined; for this purpose it does not matter how the tangent lines are determined. Thus, a variety of imaging systems and techniques can be used to capture images of an object that can be used for edge detection. In some cases, more than four tangents can be determined in a given slice. For example, more than two vantage points can be provided.

[00102] In one alternative embodiment, three cameras can be used to capture images of an object. FIG. 13 is a simplified system diagram for a system 1300 with three cameras 1302, 1304, 1306 according to an embodiment of the present invention. Each camera 1302, 1304, 1306 provides a vantage point 1308, 1310, 1312 and is oriented toward an object of interest 1313. In this embodiment, cameras 1302, 1304, 1306 are arranged such that vantage points 1308, 1310, 1312 lie in a single line 1314 in 3D space. Two-dimensional slices can be defined as described above, except that all three vantage points 1308, 1310, 1312 are included in each slice. The optical axes of cameras 1302, 1304, 1306 can be but need not be aligned, as long as the locations of vantage points 1308, 1310, 1312 are known. With three cameras, six tangents to an object can be available in a single slice. FIG. 14 illustrates a cross-section 1402 of an object as seen from vantage points 1308, 1310, 1312. Lines 1408, 1410, 1412, 1414, 1416, 1418 are tangent lines to cross-section 1402 from vantage points 1308, 1310, 1312, respectively.

[00103] For any slice with five or more tangents, the parameters of an ellipse are fully determined, and a variety of techniques can be used to fit an elliptical cross-section to the tangent lines. FIG. 15 illustrates one technique, relying on the “centerline” concept illustrated above in FIG. 9. From a first set of four tangents 1502, 1504, 1506, 1508 associated with a first pair of

5 vantage points, a first intersection region 1510 and corresponding centerline 1512 can be determined. From a second set of four tangents 1504, 1506, 1514, 1516 associated with a second pair of vantage points, a second intersection region 1518 and corresponding centerline 1520 can be determined. The ellipse of interest 1522 should be inscribed in both intersection regions. The center of ellipse 1522 is therefore the intersection point 1524 of centerlines 1512 and 1520. In this
10 example, one of the vantage points (and the corresponding two tangents 1504, 1506) are used for both sets of tangents. Given more than three vantage points, the two sets of tangents could be disjoint if desired.

[00104] Where more than five tangent points (or other points on the object's surface) are available, the elliptical cross-section is mathematically overdetermined. The extra information can
15 be used to refine the elliptical parameters, e.g., using statistical criteria for a best fit. In other embodiments, the extra information can be used to determine an ellipse for every combination of five tangents, then combine the elliptical contours in a piecewise fashion. Alternatively, the extra information can be used to weaken the assumption that the cross-section is an ellipse and allow for a more detailed contour. For example, a cubic closed curve can be fit to five or more tangents.

20 [00105] In some embodiments, data from three or more vantage points is used where available, and four-tangent techniques (e.g., as described above) can be used for areas that are within the field of view of only two of the vantage points, thereby expanding the spatial range of a motion-capture system.

[00106] While thus far the invention has been described with respect to specific embodiments,
25 one skilled in the art will recognize that numerous modifications are possible. The techniques described above can be used to reconstruct objects from as few as four tangent lines in a slice, where the tangent lines are defined between edges of a projection of the object onto a plane and two different vantage points. Thus, for purposes of the analysis techniques described herein, the edges of an object in an image are of primary significance. Any image or imaging system that supports
30 determining locations of edges of an object in an image plane can therefore be used to obtain data for the analysis described herein.

[00107] For instance, in embodiments described above, the object is projected onto an image plane using two different cameras to provide the two different vantage points, and the illuminated edge points are defined in the image plane of each camera. However, those skilled in the art with
35 access to the present disclosure will appreciate that it may be possible to use a single camera to capture motion and/or determine the shape and position of the object in 3D space.

5 [00108] One skilled in the art with access to the present disclosure will appreciate that it is possible to use more or fewer than two cameras to capture motion and/or determine the shape and position of the object in 3D space. Referring to FIG. 16A, in some embodiments, the motion-capture system 1600 includes a single camera 1602. A cross-section 1604 of the object as described above may be fit to an ellipse or any other simple closed curve. If an ellipse is used, the ellipse can be characterized by five parameters, namely, the x and y coordinates of the elliptical center (x_C , y_C), the semimajor axis (a), the semiminor axis (b), and a rotation angle (θ) (e.g., the angle of the semimajor axis relative to the x axis); five equations specify the five characteristic parameters, thereby identifying the ellipse. In various embodiments, the single camera 1602 has a vantage point 1606 that can detect two light rays 1608, 1610 transmitted from a left-edge tangent point 1612 and a right-edge tangent point 1614, respectively, on the cross-section 1604. The two tangent points 1612, 1614 define a viewed portion 1616 of the cross-section 1604 within which the portion of the cross-section is within the field of view of the camera. The two tangent points 1612, 1614 provide two equations that can partially determine the ellipse 1618 that fits most closely to the cross-section 1604.

20 [00109] FIG. 16B illustrates a motion-capture system 1600 that includes three light sources (e.g., LEDs) 1620, 1622, 1624 to illuminate the object and provide additional information about the cross-section 1604 to determine the parameters of the ellipse. In one embodiment, depending on the relative positions between the object, the camera 1602, and the light sources 1620, 1622, 1624, the camera 1602 detects shadow (e.g., unilluminated) regions on the cross-section 1604 created by the light sources 1620, 1622, 1624. For example, the light source 1620 illuminates a partial portion 1626 of the cross-section 1604; the illuminated portion 1626 is determined by the position of the light source 1620 and two illuminated edge points 1628, 1630. For example, the two illuminated edge points 1628, 1630 may be defined by the light rays 1632, 1634 tangent to the cross-section 1604. As a result, a shadow (or unilluminated) region 1636 that has limited exposure to the light illumination from the light source 1620 can be observed on the cross-section 1604.

30 [00110] Because the viewed portion of the cross-section 1604 within the camera's field of view is defined by the two tangent points 1612, 1614, the camera 1602 can detect the illuminated part 1638 between the tangent points 1612 and 1628 and the shadow region 1640 between the tangent points 1628 and 1614. Accordingly, the boundary between the shadow region 1640 and the illuminated part 1638 defines a shadow edge point 1628. In various embodiments, the shadow edge point 1628 is detected by the camera 1602; the detected shadow point 1628 can provide two additional equations that further determine the characteristic ellipse parameters. The first equation

5 is based on the detected position of the shadow edge point 1628, and the second equation is based on a light ray 1642 emitted from the light source 1620 to the shadow edge point 1628. In one embodiment, the path of the light ray 1642 is based on the spatial relationship between the camera 1602 and the light source 1616. Although the detected shadow edge point 1628 provides two additional equations to determine the characteristic parameters of the ellipse, the shadow edge point 10 1628 introduces an additional unknown parameter (e.g. the distance between the camera 1602 and the shadow edge point 1628).

[00111] FIG. 16C shows how multiple light sources 1620, 1622, 1624 illuminate the cross-section 1604 and generate multiple shadow regions 1644, 1646, 1648, respectively. Shadow edge points 1628, 1650, 1652 defining boundaries between the shadow regions 1644, 1646, 1648 and the 15 illuminated regions are detected by the camera 1602. As described above, each detected shadow edge point 1628, 1650, 1652 can provide two additional equations that contribute to determining the characteristic ellipse parameters. In addition, each detected shadow edge point 1628, 1650, 1652 may also introduce an additional unknown parameter (e.g. the distance between the camera 1602 and the corresponding shadow edge point). As a result, utilizing three light sources 1620, 20 1622, 1624 creates three additional unknown parameters and provides six equations to determine the characteristic the ellipse parameters.

[00112] In summary, when the motion-capture system 1600 includes three light sources 1620, 1622, 1624, each creating a shadow edge point on the cross-section 1604 of the object, eight unknown parameters — including the five characteristic parameters of the ellipse and the three 25 distances from the camera 1602 to the three shadow edge points 1628, 1646, 1656 — are solved to determine the position, rotation, and size of the ellipse. In various embodiments, the camera 1602 detects two tangent points 1612, 1614 on the object cross-section that provide two equations to partially determine the eight unknown parameters and the three shadow edge points 1628, 1646, 1656 created by light emitted from the light sources 1620, 1622, 1624 onto the cross-section 1604 30 of the object. Because each shadow edge point 1628, 1646, 1656 provides two equations, the motion-capture system 1600 thus has eight equations to solve for the eight unknown parameters, including five unknown ellipse parameters and the three unknown distances introduced by the three shadow edge points 1628, 1646, 1656.

[00113] In some embodiments, the motion-capture system 1600 includes fewer than three light 35 sources, such as two light sources 1620, 1622. The two light sources 1620, 1622 create two unknown parameters (i.e., the distance between the camera and the shadow edge points generated by the light sources 1620, 1622) and four equations (i.e., the two positions of the shadow edge

5 points and the two light rays emitted from the light sources 1620, 1622 to the shadow edge points). Accordingly, the motion-capture system 1600 has, in total, seven unknown parameters and six equations, so the ellipse is underdetermined. Because the six equations by themselves cannot have an exact solution for the seven unknown parameters, in one embodiment, one of the seven unknown parameters is initially estimated. The estimated parameter is then applied to the six equations, and
10 the other six unknown parameters can be solved. In one embodiment, the self-consistency of the estimated parameter and the six solved parameters is checked in the end of the process to determine the accuracy of the estimated parameter. This process may iterate until a maximum self-consistency or accuracy of the estimated parameter is obtained.

[00114] In some embodiments, there are more than three light sources, yielding more available
15 equations than unknown parameters, and the ellipse is overdetermined. The extra equations may be utilized to refine the elliptical parameters, e.g., using statistical criteria for a best fit. Alternatively, the extra information can be used to weaken the assumption that the cross-section is an ellipse and allow for a more detailed contour. For example, a cubic closed curve can be fit to the two tangent points and the three shadow points. In some embodiments, the extra equations may be used to
20 optimize the speed and/or accuracy of a numerical solver that is utilized to solve the unknown parameters and is implemented on a general-purpose computing device.

[00115] In some embodiments, the light sources 1620, 1622, 1624 are pulsed on individually and in succession, allowing for each of the shadow regions 1644, 1646, 1648 and the shadow points
25 1628, 1650, 1652 to be associated with a single light source. In some embodiments, light sources 1616, 1618, and 1620 may be of different wavelengths, allowing the shadow edge points 1628, 1650, 1652 to be easily identified.

[00116] Additionally, those skilled in the art with access to the present disclosure will appreciate that cameras are not the only tool capable of projecting an object onto an imaging surface. For example, a light source can create a shadow of an object on a target surface, and the shadow —
30 captured as an image of the target surface — can provide a projection of the object that suffices for detecting edges and defining tangent lines. The light source can produce light in any visible or non-visible portion of the electromagnetic spectrum. Any frequency (or range of frequencies) can be used, provided that the object of interest is opaque to such frequencies while the ambient environment in which the object moves is not. The light sources used should be bright enough to
35 cast distinct shadows on the target surface. Point-like light sources provide sharper edges than diffuse light sources, but any type of light source can be used.

5 [00117] In one such embodiment, a single camera is used to capture images of shadows cast by multiple light sources. FIG. 17 illustrates a system 1700 for capturing shadows of an object according to an embodiment of the present invention. Light sources 1702 and 1704 illuminate an object 1706, casting shadows 1708, 1710 onto a front side 1712 of a surface 1714. Surface 1714 can be translucent so that the shadows are also visible on its back side 1716. A camera 1718 can be oriented toward back side 1716 as shown and can capture images of shadows 1708, 1710. With this arrangement, object 1706 does not occlude the shadows captured by camera 1718. Light sources 1702 and 1704 define two vantage points, from which tangent lines 1720, 1722, 1724, 1726 can be determined based on the edges of shadows 1708, 1710. These four tangents can be analyzed using techniques described above.

15 [00118] In an embodiment such as system 1700 of FIG. 17, shadows created by different light sources may partially overlap, depending on where the object is placed relative to the light source. In such a case, an image may have shadows with penumbra regions (where only one light source is contributing to the shadow) and an umbra region (where the shadows from both light sources overlap). Detecting edges can include detecting the transition from penumbra to umbra region (or vice versa) and inferring a shadow edge at that location. Since an umbra region will be darker than a penumbra region; contrast-based analysis can be used to detect these transitions.

[00119] Certain physical or object configurations may present ambiguities that are resolved in accordance with various embodiments we as now discussed. Referring to FIG. 18, when two objects 1808, 1810 are present, the camera 1820 may detect four shadows 1812, 1814, 1816, 1818 and the tangent lines may create four intersection regions 1822, 1824, 1826, 1828 that all lie within the shadow regions 1830, 1832, 1834, 1836. Because it is difficult to determine, from a single slice of the shadow image, which of these intersection regions contain portions of the object, an analysis of whether the intersection regions 1822, 1824, 1826, 1828 are occupied by the objects may be ambiguous. For example, shadows 1812, 1814, 1816, 1818 that are generated when intersection regions 1822 and 1826 are occupied are the same as those generated when regions 1824 and 1828 are occupied, or when all four intersection regions 1822, 1824, 1826, 1828 are occupied. In one embodiment, correlations across slices are used to resolve the ambiguity in interpreting the intersection regions (or “visual hulls”) 1822, 1824, 1826, 1828.

35 [00120] In various embodiments, referring to FIG. 19, a system 1900 incorporates a large number of light sources (i.e., more than two light sources) to resolve the ambiguity of the intersection regions when there are multiple objects casting shadows. For example, the system 1900 includes three light sources 1902, 1904, 1906 to cast light onto a translucent surface 1910 and

a camera 1912 positioned on the opposite side of surface 1910 to avoid occluding the shadows cast by an object 1914. As shown in FIG. 19, because utilization of three light sources provides five or more tangents for one or more objects 1914 in a slice, the ellipse-fitting techniques described above may be used to determine the cross-sections of the objects. A collection of the cross-sections of the objects in 2D slices may then determine the locations and/or movement of the objects.

[00121] If multiple objects, however, are located in close proximity (e.g., the fingers of a hand), utilization of additional light sources may reduce the sizes of the various intersection regions as well as increase the total number of intersection regions. If the number of light sources is much greater than the number of the proximal objects, the intersection regions may be too small to be analyzed based on a known or assumed size scale of the object. Additionally, the increased number of intersection regions may result in more ambiguity in distinguishing intersection regions that contain objects from intersection regions that do not contain objects (i.e., “blind spots”). In various embodiments, whether an intersection region contains an object is determined based on the properties of a collection of intersection points therein. As described in greater detail below, an intersection point is defined by at least two shadow lines, each connecting a shadow point of the shadow and a light source. If the intersection points in an intersection region satisfy certain criteria, the intersection region is considered to have the objects therein. A collection of the intersection regions may then be utilized to determine the shape and movement of the objects.

[00122] Referring to FIG. 20, a collection of the intersection regions (a visual hull) 2030 is defined by a virtual rubber band 2032 stretched around multiple intersection regions 2031 (or “convex hulls”); each intersection region 2031 is defined by a smallest set of intersection points 2034. When there are multiple intersection regions 2031, distinguishing each intersection region 2031 from a collection of intersection points 2034 may be difficult. In some embodiments, referring to FIG. 21, a simple visual hull is first constructed by a setup of two lights 2102, 2104 (here denoted L_n , with $n=\{1, 2\}$ to permit further generalization to greater numbers of light

sources, shadows, shadow regions, points, and visual hulls), each casting one shadow 2106A, 2106B, respectively. The light source L_1 and shadow 2106A define a shadow region, $R_{1,1}$; similarly, light source L_2 and the shadow 2106B define a shadow region, $R_{2,1}$; in general, the shadow region is denoted as, $R_{u,v}$, where u is the number of the corresponding light source and v is a number that denotes a left to right ordering in a scene within the set of all shadow regions from the light source u . Boundaries of the shadows (or “shadow points”) lie on an x axis and are denoted by $S_{u,v}$. The shadow points and each light source may then create shadow lines 2108, 2110, 2112, 2114; the shadow lines are referenced by the two connecting points; for example, $\overline{L_1 S_{1,2}}$,

5 (abbreviated $\overline{S_{1,2}}$, where the first subscript also refers to the light number). The convex hull 2130 (or visual hull here since there is only one intersection region 2128) may then be defined by the four intersection points 2134 in the example of FIG. 21. In one embodiment, the intersection points 2134 are determined based on the intersections of every pair of shadow lines, for example, $\overline{S_{1,1}}$, $\overline{S_{1,2}}$, $\overline{S_{2,1}}$, and $\overline{S_{2,2}}$. Because pairs of shadow lines from the same light source L_1 or L_2 do not

10 intersect, the intersection of the pairs of lines from the same light source may then be neglected. [00123] When there are more than two light sources, determining all shadow line intersections no longer suffices to find intersection points that lie on the intersection region 2128. Referring to FIG. 22A, utilization of three light sources 2202, 2204, 2206, may result in “true” intersection points 2234A, 2234B, 2234C, 2234D, 2234E, 2234F that form the intersection region 2228

15 occupied by the object 2208 and “false” intersection points 2235A, 2235B, 2235C, 2235D, 2235E, 2235F that clearly do not form the intersection region 2228. For example, the false intersection point 2235E created by a left shadow line 2224 of the shadow region 2218A and a right shadow line 2226 of the shadow region 2218B is a false intersection point because it does not lie inside the intersection region 2228. Because the intersection region 2228 is an intersection of the shadow

20 regions 2218A, 2218B, 2218C created by the object 2208 and the light sources 2202, 2204, 2206, the number of shadow regions in which each “true” intersection point lies is equal to the number of the light sources (i.e., three in FIG. 22A). “False” intersection points, by contrast, lie outside the intersection region 2228 even though they may lie inside an intersection region that includes fewer number of shadow regions compared to the total number of light sources. In one embodiment,

25 whether an intersection point is “true” or “false” is determined based on the number of shadow regions included in the intersection region in which the intersection point lies. For example, in the presence of three light sources in FIG. 22A, the intersection point 2234A is a true intersection point because it lies inside three shadow regions 2218A, 2218B, 2218C; whereas the intersection point 2235F is a false intersection point because it lies inside only two shadow regions 2218B, 2218C.

30 [00124] Because the intersection regions are defined by a collection of intersection points, excessive computational effort may be required to determine whether an intersection point is contained by a correct number of regions (i.e., the number of the light sources). In some embodiments, this computational complexity is reduced by assuming that each intersection point is not “false” and then determining whether the results are consistent with all of the shadows captured

35 by the camera. These configurations project each intersection point $I = [I_x, I_y]$ onto the x axis through a ray directed from each light source $L = [L_x, L_y]$ that is not involved in the original

5 intersection determination. The solutions for these projections are given by $\left[\frac{L_y P_x - L_x P_y}{L_y - P_y}, 0 \right]$. If a projection point on the x axis lies inside a shadow region from the testing light source, it is likely that the projected intersection point is a true intersection point. For example, referring to FIG. 22B, the intersection point 2235E is determined by the shadow lines 2224 and 2226 created by the light sources 2202 and 2206. Projecting the intersection point 2235E onto the x axis using the light

10 source 2206, which is not involved in determining the intersection point 2235E, creates a projection point P_3 . Because the projection point P_3 does not lie inside the shadow region 2218C created by the light source 2206 and the object 2208, the intersection point 2235E is considered to be a false intersection point; whereas the intersection point 2234E is a true intersection point because the projection point P_1 thereof lies within the shadow region 2218A. As a result, for every possible

15 intersection point, an additional $N - 2$ projections must be determined for the $N - 2$ light sources that are not involved in determining the position of the intersection point (where N is the total number of light sources in the system). In other words, a projection check must be made for every light source other than the original two that are used to determine the tested intersection point. Because determining whether the intersection point is true or false based on the projections is

20 simpler than checking the number of shadow regions in which each intersection point lies, the required computational requirements and processing time may be significantly reduced.

[00125] If, however, a large quantity of light sources is utilized in the system, the overall process may still be time-consuming. In various embodiments, the light sources L_1 , L_2 , and L_3 are placed in a line parallel to the x axis, the location of the projection points can then be determined without

25 finding the location of the intersection point for every pair of shadow lines. Accordingly, whether the intersection point 2234 is a true or false point may be determined without finding or locating the position thereof; this further reduces the processing time. For example, with reference to FIG. 22C, assuming that the shadow points S_1 and S_3 are either known or have been determined, whether the intersection point I of the shadow lines $\overline{L_1 S_3}$ and $\overline{L_3 S_1}$ is true or false may be determined by the

30 position of the projection point P_2 created by the light source L_2 . The distance between the projection point P_2 created by the light source L_2 and the shadow point S_1 is given as:

$$\overline{S_1 P_2} = \overline{S_1 S_3} \left[\frac{\overline{L_2 L_3}}{\overline{L_1 L_3}} \right] \quad (\text{Eq. 1})$$

Thus, the location of any one of the projection points projected from the intersection point, I , and light sources may be determined based on the other two shadow points and the distance ratios

5 associated with light sources L_1 , L_2 and L_3 . Because the ratio of the distances between the light sources is predetermined, the complexity in determining the projection point P_2 is reduced to little more than calculating distances between the shadow points and multiplying these distances by the predetermined ratio. If the distance between the projection point P_2 and the shadow point S_1 is larger than the size of the shadow, i.e., $\overline{S_1 S_3}$, that is captured by the camera, the intersection point, 10 I , is a false point. If, on the other hand, the distance between the projection points S_2 and S_1 is smaller than the size of the shadow, the intersection point I is likely a true point. Although the location of the intersection point, I , may still be determined based on the shadow lines $\overline{L_1 S_3}$ and $\overline{L_3 S_1}$, this determination may be skipped during the process. Accordingly, by aligning the light sources in a line, the false intersection points can be quickly determined without performing the 15 complex computations, thereby saving a large amount of processing time and power.

[00126] More generally, when there are N light sources, each denoted as L_i ($1 \leq i \leq N$), arranged on a line parallel to the x axis and each light source possesses a set of S_i shadow points (where i is the light number), a total number of M intersection calculations for all possible intersection pairs is given as:

$$20 \quad M = \sum_{i=1}^{N-1} S_i \left(\sum_{k=i+1}^N S_k \right). \quad (\text{Eq. 2})$$

For example, if there are N light sources, each casting n shadows, the total number of intersection calculations M may then be given as

$$M = n^2 N(N-1). \quad (\text{Eq. 3})$$

Because each of these intersection calculations involves multiple operations (e.g., addition and 25 multiplication), the total number of operations, T_o , may be given as

$$T_o = 2n^2 N(2N+1)(N-1). \quad (\text{Eq. 4})$$

For example, a total number of operations $T_o = 2(1)^2 3(2 \bullet 3 + 1)(3 - 1) = 84$ is required to determine the simplest visual hull 2128 shown in FIG. 21. In one embodiment, there are, for example, 12 light sources (i.e., $N = 12$), each casting 10 shadows (i.e., $n = 10$); the number of required 30 intersection calculations for this scenario is $M = 13,200$, setting the number of total operations to be $T_o = 660,000$. Again, this requires a significant amount of processing time. In some embodiments, the distance ratios between light sources are predetermined, and as a result, only one operation (i.e., multiplication) is needed to determine which pairs of shadow points produce true intersection points; this reduces the number of total operations to 13,200.

5 [00127] The computational load required to find the visual hull depends on the quantity of the true intersection points, which may not be uniquely determined by the number of shadows. Suppose, for example, that there are N light sources and each object is a circle that casts one shadow per light; this results in N intersection regions (or $6N$ intersection points) per object. Because there are n objects, the resulting number of intersection points that need to be checked is
 10 $6Nn^2$ (i.e., roughly 6,000 for 10 objects cast by 12 light sources). As described above, the number of operations required for the projection check is 13,200; accordingly, a total number of operations 19,200 is necessary to determine the visual hull formed by the true intersection points. This is a 34-fold improvement in determining the solution for a single 2D scene compared to the previous estimate of 660,000 operations. The number of reduced operations may be given as:

$$15 \quad T_p = n^2 N(N-1) + 6Nn^2 \quad (\text{Eq. 5})$$

The ratio of the required operations to the reduced operations may then be expressed as:

$$\frac{T_o}{T_p} = \frac{2n(2N+1)(N-1)}{nN - n + 6n} \quad (\text{Eq. 6})$$

Based on Eq. 6, if the light sources lie along a line or lines parallel to the x axis, the improvement is around an order of magnitude for a small number of lights, whereas the improvement is nearly two
 20 orders of magnitude for a larger number of lights.

[00128] If the objects are reconstructed in 3D space and/or a fast real-time refresh rate (e.g., 30 frames per second) is used by the camera, the computational load may be increased by several orders of magnitude due to the additional complexity. In some embodiments, the visual hull is split into a number of small intersection regions that can generate at least a portion of the shadows in the
 25 scene; the smallest cardinality of the set of small intersection regions is defined as a “minimal solution.” In one embodiment, the number of the small intersection regions in the minimal solution is equal to the largest number of shadows generated by any single light source. The computational complexity of obtaining the visual hull may significantly be reduced by determining each of the small visual hulls prior to assembling them together into the visual hull.

30 [00129] Referring again to FIG. 20, the intersection points 2034 may form an amorphous cloud that does not imply particular regions. In various embodiments, this cloud is first split into a number of sets, each set determining an associated convex hull 2031. As further described below, in one embodiment, a measure is utilized to determine the intersection region to which each intersection point belongs. The determined intersection region may then be assembled into an exact
 35 visual hull. In one implementation, the trivial case of a visual hull containing only one intersection region is ignored. In some embodiments, every intersection region ρ is assigned an N -dimensional

- 5 subscript, where N is the number of light sources in the scene under consideration. The n th entry for this subscript of the intersection region ρ is defined as the value v of the u th subscript (where $u = n$) for each shadow region $R_{u,v}$ of which the intersection region is a subset; every intersection region thus has a unique identifier for grouping the intersection points, as shown in FIG. 23. Because two of the subscript entries for an intersection point can be determined directly from the
- 10 two shadow lines, the resulting intersection point thereof is in the two shadow regions in which the shadow lines are located. For the rest of the entries, the locations of the projections of the intersection points may be recorded during the determination of true and false intersection points. Complete knowledge of the particular intersection regions for each intersection point may thus be determined.
- 15 [00130] Once the distinct intersection regions have been determined, the smallest subset of intersection regions that can generate all of the final shadows may then be found. FIG. 23 depicts intersection regions $\rho_{1,1,1}$, $\rho_{2,2,2}$, $\rho_{3,3,3}$ resulting from casting light from three light sources onto three objects 2338A, 2338B, and 2338C. Because the greatest number of shadows cast by any particular light source in this case is three and the number of intersection regions in the minimal
- 20 solution is equal to the largest number of shadows generated by any single light source, every group that includes three intersection regions in the scene may be tested. If a group generates a complete set of shadows captured by the camera, this group is the minimum solution. The number of trios to test is equal to the binomial coefficient $C_u^j = \binom{j}{u} = \frac{j!}{u!(j-u)!}$, where j is the total number of
- intersection regions. For example, there are $C_3^{13} = 286$ combinations in FIG. 23. The likelihood
- 25 that a trio having larger intersection regions can generate all of the captured shadows is higher than for a trio having smaller intersection regions; additionally, larger intersection regions usually have a greater number of intersection points. In some embodiments, the number of trios tested is reduced by setting a criterion value U equal to the greatest number of intersection points in any intersection region. For example, only regions or combinations of regions having a number of intersection
- 30 points exceeding the criteria number U are checked. If there are no solutions, U may be reset to $U - 1$ and the process may be repeated. For example, by setting $U = 6$, there are only five regions, $\rho_{1,1,1}$, $\rho_{2,2,2}$, $\rho_{3,3,3}$, $\rho_{1,2,3}$, $\rho_{3,2,1}$, having six intersection points need to be checked. The region subscripts may be presented as a single number vector, e.g., $\rho_{1,1,1} = [1 \ 1 \ 1]$; and the combination

5 of $\rho_{3,2,1}$, $\rho_{1,1,1}$, and $\rho_{2,2,2}$ may be written as a matrix, e.g., $\begin{bmatrix} 3 & 2 & 1 \\ 1 & 1 & 1 \\ 2 & 2 & 2 \end{bmatrix}$. There are nine additional

combinations exist in FIG. 23:

$$\begin{bmatrix} 3 & 2 & 1 \\ 1 & 1 & 1 \\ 3 & 3 & 3 \end{bmatrix}, \begin{bmatrix} 3 & 2 & 1 \\ 1 & 1 & 1 \\ 1 & 2 & 3 \end{bmatrix}, \begin{bmatrix} 3 & 2 & 1 \\ 2 & 2 & 2 \\ 3 & 3 & 3 \end{bmatrix}, \begin{bmatrix} 3 & 2 & 1 \\ 2 & 2 & 2 \\ 1 & 2 & 3 \end{bmatrix}, \begin{bmatrix} 3 & 2 & 1 \\ 3 & 3 & 3 \\ 1 & 2 & 3 \end{bmatrix}, \begin{bmatrix} 1 & 1 & 1 \\ 2 & 2 & 2 \\ 3 & 3 & 3 \end{bmatrix}, \begin{bmatrix} 1 & 1 & 1 \\ 2 & 2 & 2 \\ 1 & 2 & 3 \end{bmatrix},$$

$$\begin{bmatrix} 1 & 1 & 1 \\ 3 & 3 & 3 \\ 1 & 2 & 3 \end{bmatrix}, \begin{bmatrix} 2 & 2 & 2 \\ 3 & 3 & 3 \\ 1 & 2 & 3 \end{bmatrix}.$$

10 Because the minimal solution alone can generate all of the shadows in the scene, each column of the minimal solution matrix has the numbers 1, 2, 3 (in no particular order). Accordingly, the 6th combination above having $\rho_{1,1,1}$, $\rho_{2,2,2}$, and $\rho_{3,3,3}$ is the minimal solution. This approach finds the minimal solution by determining whether there is at least one intersection region in every shadow region. This approach, however, may be time-consuming upon reducing U to 3, as the regions that have three intersection point require a more complicated check. In some embodiments, the three-

15 point regions are neglected since they are almost never a part of a minimal solution.

[00131] In some embodiments, the 3D scenes are decomposed into a number of 2D scenes that can be quickly solved by the approaches as described above to determine the 3D shape of the objects. Because many of these 2D scenes share the same properties (e.g., the shape or location of the intersection regions), the solution of one 2D slice may be used to determine the solution of the

20 next 2D slice; this may improve the computational efficiency.

[00132] The light sources may be positioned to lie in a plane. In one embodiment, a number of “bar” light sources are combined with “point” light sources to accomplish more complex lighting arrangements. In another embodiment, multiple light arrays lying in a plane are combined with multiple outlier-resistant least squares fits to effectively reduce the computational complexity by

25 incorporating previously known geometric parameters of the target object.

[00133] Referring to FIG. 24, in some embodiments, a shadow 2412 is cast on a translucent or imaginary surface 2440 such that the shadow 2412 can be viewed and captured by a camera 2438. The camera 2438 may take pictures with a number of light sensors (not shown in FIG. 24) arranged in a rectangular grid. In the camera 2438, there may be three such grids interlaced at small

30 distances that essentially lie directly on top of each other. Each grid has a different color filter on all of its light sensors (e.g., red, green, or blue). Together, these sensors output three images, each

5 comprising $A \times B$ light brightness values in the form of a matrix of pixels. The three color images together form an $A \times B \times 3$ RGB image matrix. The image matrices may have their own coordinate system, which is defined by the set of matrix cell subscripts for a given pixel. For example, indices $(x, y, z) = (0, 0, 0)$ may be defined and start in an upper left corner 2439 of the image. In one embodiment, the matrix of $z = 1$ represents the red color image and $z = 2$ and $z = 3$ are the green and blue images, respectively. In one implementation, an “image row” is defined as all pixel values for a given constant coordinate value of y and an “image column” is defined as all pixel values for a given constant coordinate value of x .

[00134] Referring to FIG. 25A, a color image 2550 is split into images 2552, 2554, 2556 of three primary colors (i.e., red, green, and blue, respectively) by decomposing an $A \times B \times 3$ full color matrix in a memory into 3 different $A \times B$ matrices, one for each z value between 1 and 3. Pixels in each image 2552, 2554, 2556 are then compared to a brightness threshold value to determine which pixels represent shadow and which represent background to thereby generate three shadow images 2558, 2560, 2562, respectively. The brightness threshold value may be determined by a number of statistical techniques. For example, in some embodiments, a mean pixel brightness is determined for each image and the threshold is set by subtracting three times the standard deviation of the brightness of the same pixels in the same image. Edges of the shadow images 2558, 2560, 2562 may then be determined to generate shadow point images 2564, 2566, 2568, respectively, using a conventional edge-determining technique. For example, the edge of each shadow image may be determined by subtracting the shadow image itself from an offset image created by offsetting a single pixel on the left (or right, top and/or bottom) side thereof. The 2D approaches described above may be applied to each of the shadow point images 2564, 2566, 2568 to determine the locations and colors of the objects. In some embodiments, shadow points in images 2564, 2566, 2568 are combined into a single $A \times B \times 3$ color matrix or image 2570. Application of the 2D approaches described above to the combined shadow point image 2570 can then reconstruct an image of the object 2572 (e.g., a hand, as shown in FIG. 25B). Reconstructing an object (e.g., a hand) from shadows using various embodiments in the present invention may then be as simple as reconstructing a number of 2D ellipses. For example, fingers may be approximated by circles in 2D slices, and a palm may be approximated as an ellipse. This reconstruction is thereby converted into a practical number of simpler, more efficient reconstructions; the reconstructed 2D slices are then reassembled into the final 3D solution. These efficient reconstructions may be computed using a single processor or multiple processors operating in parallel to reduce the processing time.

[00135] In various embodiments, referring again to FIG. 24, the image coordinate system (i.e.,

5 the “imaging grid” 2442) is imposed on the surface 2440 to form a standard Cartesian coordinate system thereon such that the shadow 2412 can be easily defined. For example, each pixel (or light measurement value) in an image may be defined based on the coordinate integers x and y . In some embodiments, the camera 2438 is perpendicular to the surface 2440 on which shadows 2412 are cast and a point on a surface in the image grid is defined based on its coordinate inside an image
10 taken by the camera 2438. In one embodiment, all light sources lie along a line or lines on a plane perpendicular to one of the axes to reduce the computational complexity. In various embodiments, the z axis of the coordinate system uses the same distance units and is perpendicular to the x and y axes of image grid 2442 to capture the 3D images of the shadows. For example, the light sources may be placed parallel to the x or y axis and perpendicular to the z -axis; a 3D captured shadow
15 structure in the image coordinate system may be split into multiple 2D image slices, where each slice is a plane defined by a given row on the imaging grid and the line of light sources. The 2D slices may or may not share similar shapes. For example, the 2D intersection region of a 3D intersection region for a spherical object is very similar, i.e., a circle; whereas the 2D intersection region of a 3D intersection region for a cone shape varies across the positions of the 2D slices.

20 [00136] As described above, the shape of multiple objects may be discerned by determining a minimal solution of each 2D slice obtained from the 3D shadow. Since two slices next to each other are typically very similar, multiple slices often have the same minimal solution. In various embodiments, when two nearby slices have the same number of intersection regions, different combinations of the intersection regions are bypassed between the slices and the combination that
25 works for a previous slice is reused on the next slice. If the old combination works for the new slice, this solution becomes a new minimal solution for the new slice and any further combinatorial checks are not performed. The reuse of old combinations thus greatly reduces computational time and complexity for complicated scenes. Although various embodiments described above are related to determining the shapes and positions of objects in 3D space using cross-sections obtained
30 from the shadows cast by the objects, one of ordinary skill in the art will understand that cross-sections obtained utilizing different approaches, e.g., reflections from the objects, are within the scope of the current invention.

[00137] In still other embodiments, a single camera can be used to capture an image of both the object and one or more shadows cast by the object from one or more light sources at known
35 positions. Such a system is illustrated in FIGs. 26A and 26B. FIG. 26A illustrates a system 2600 for capturing a single image of an object 2602 and its shadow 2604 on a surface 2606 according to an embodiment of the present invention. System 2600 includes a camera 2608 and a light source 2612

5 at a known position relative to camera 2608. Camera 2608 is positioned such that object of interest 2602 and surface 2606 are both within its field of view. Light source 2612 is positioned so that an object 2602 in the field of view of camera 2608 will cast a shadow onto surface 2606. FIG. 26B illustrates an image 2620 captured by camera 2608. Image 2620 includes an image 2622 of object 2602 and an image 2624 of shadow 2604. In some embodiments, in addition to creating shadow
10 2604, light source 2612 brightly illuminates object 2602. Thus, image 2620 will include brighter-than-average pixels 2622, which can be associated with illuminated object 2602, and darker-than-average pixels 2624, which can be associated with shadow 2604.

[00138] In some embodiments, part of the shadow edge may be occluded by the object. Where
15 the object can be reconstructed with fewer than four tangents (e.g., using circular cross-sections), such occlusion is not a problem. In some embodiments, occlusion can be minimized or eliminated by placing the light source so that the shadow is projected in a different direction and using a camera with a wide field of view to capture both the object and the unoccluded shadow. For example, in FIG. 26A, the light source could be placed at position 2612'.

[00139] In other embodiments, multiple light sources can be used to provide additional visible
20 edge points that can be used to define tangents. For example, FIG. 26C illustrates a system 2630 with a camera 2632 and two light sources 2634, 2636, one on either side of camera 2632. Light source 2634 casts a shadow 2638, and light source 2636 casts a shadow 2640. In an image captured by camera 2632, object 2602 may partially occlude each of shadows 2638 and 2640. However, edge 2642 of shadow 2638 and edge 2644 of shadow 2640 can both be detected, as can the edges of
25 object 2602. These points provide four tangents to the object, two from the vantage point of camera 2632 and one each from the vantage point of light sources 2634 and 2636.

[00140] As yet another example, multiple images of an object from different vantage points can be generated within an optical system, e.g., using beamsplitters and mirrors. FIG. 27 illustrates an image-capture setup 2700 for a motion capture system according to another embodiment of the
30 present invention. A fully reflective front-surface mirror 2702 is provided as a "ground plane." A beamsplitter 2704 (e.g., a 50/50 or 70/30 beamsplitter) is placed in front of mirror 2702 at about a 20-degree angle to the ground plane. A camera 2706 is oriented toward beamsplitter 2704. Due to the multiple reflections from different light paths, the image captured by the camera can include ghost silhouettes of the object from multiple perspectives. This is illustrated using representative
35 rays. Rays 2706a, 2706b indicate the field of view of a first virtual camera 2708; rays 2710a, 2710b indicate a second virtual camera 2712; and rays 2714a, 2714b indicate a third virtual camera 2716.

5 Each virtual camera 2708, 2712, 2716 defines a vantage point for the purpose of projecting tangent lines to an object 2718.

[00141] Another embodiment uses a screen with pinholes arranged in front of a single camera.

FIG. 28 illustrates an image capture setup 2800 using pinholes according to an embodiment of the present invention. A camera sensor 2802 is oriented toward an opaque screen 2804 in which are
10 formed two pinholes 2806, 2808. An object of interest 2810 is located in the space on the opposite side of screen 2804 from camera sensor 2802. Pinholes 2806, 2808 can act as lenses, providing two effective vantage points for images of object 2810. A single camera sensor 2802 can capture images from both vantage points.

[00142] More generally, any number of images of the object and/or shadows cast by the object
15 can be used to provide image data for analysis using techniques described herein, as long as different images or shadows can be ascribed to different (known) vantage points. Those skilled in the art will appreciate that any combination of cameras, beamsplitters, pinholes, and other optical devices can be used to capture images of an object and/or shadows cast by the object due to a light source at a known position.

20 [00143] Further, while the embodiments described above use light as the medium to detect edges of an object, other media can be used. For example, many objects cast a “sonic” shadow, either blocking or altering sound waves that impinge upon them. Such sonic shadows can also be used to locate edges of an object. (The sound waves need not be audible to humans; for example, ultrasound can be used.) The term “shadow” is herein used broadly to connote light or sonic
25 shadows or other occlusion of a disturbance by an object, and the term “light” means electromagnetic radiation of any suitable wavelength(s) or wavelength range.

[00144] As described above, the general equation of an ellipse includes five parameters; where only four tangents are available, the ellipse is underdetermined, and the analysis proceeds by assuming a value for one of the five parameters. Which parameter is assumed is a matter of design
30 choice, and the optimum choice may depend on the type of object being modeled. It has been found that in the case where the object is a human hand, assuming a value for the semimajor axis is effective. For other types of objects, other parameters may be preferred.

[00145] Further, while some embodiments described herein use ellipses to model the cross-
sections, other shapes can be substituted. For instance, like an ellipse, a rectangle can be
35 characterized by five parameters, and the techniques described above can be applied to generate rectangular cross-sections in some or all slices. More generally, any simple closed curve can be fit to a set of tangents in a slice. (The term “simple closed curve” is used in its mathematical sense

5 throughout this disclosure and refers generally to a closed curve that does not intersect itself with no limitations implied as to other properties of the shape, such as the number of straight edge sections and/or vertices, which can be zero or more as desired.) The number of free parameters can be limited based on the number of available tangents. In another embodiment, a closed intersection region (a region fully bounded by tangent lines) can be used as the cross-section, without fitting a
10 curve to the region. While this may be less accurate than ellipses or other curves, e.g., it can be useful in situations where high accuracy is not desired. For example, in the case of capturing motion of a hand, if the motion of the fingertips is of primary interest, cross-sections corresponding to the palm of the hand can be modeled as the intersection regions while fingers are modeled by fitting ellipses to the intersection regions.

15 [00146] In some embodiments, cross-slice correlations can be used to model all or part of the object using 3D surfaces, such as ellipsoids or other quadratic surfaces. For example, elliptical (or other) cross-sections from several adjacent slices can be used to define an ellipsoidal object that best fits the ellipses. Alternatively, ellipsoids or other surfaces can be determined directly from tangent lines in multiple slices from the same set of images. The general equation of an ellipsoid
20 includes nine free parameters; using nine (or more) tangents from two or three (or more) slices, an ellipsoid can be fit to the tangents. Ellipsoids can be useful, e.g., for refining a model of fingertip (or thumb) position; the ellipsoid can roughly correspond to the last segment at the tip of a finger (or thumb). In other embodiments, each segment of a finger can be modeled as an ellipsoid. Other quadratic surfaces, such as hyperboloids or cylinders, can also be used to model an object or a
25 portion thereof.

[00147] In some embodiments, an object can be reconstructed without tangent lines. For example, given a sufficiently sensitive time-of-flight camera, it would be possible to directly detect the difference in distances between various points on the near surface of a finger (or other curved object). In this case, a number of points on the surface (not limited to edge points) can be
30 determined directly from the time-of-flight data, and an ellipse (or other shape) can be fit to the points within a particular image slice. Time-of-flight data can also be combined with tangent-line information to provide a more detailed model of an object's shape.

[00148] Any type of object can be the subject of motion capture using these techniques, and various aspects of the implementation can be optimized for a particular object. For example, the
35 type and positions of cameras and/or light sources can be optimized based on the size of the object whose motion is to be captured and/or the space in which motion is to be captured. As described above, in some embodiments, an object type can be determined based on the 3D model, and the

5 determined object type can be used to add type-based constraints in subsequent phases of the analysis. In other embodiments, the motion capture algorithm can be optimized for a particular type of object, and assumptions or constraints pertaining to that object type (e.g., constraints on the number and relative position of fingers and palm of a hand) can be built into the analysis algorithm. This can improve the quality of the reconstruction for objects of that type, although it may degrade
10 performance if an unexpected object type is presented. Depending on implementation, this may be an acceptable design choice. For example, in a system for controlling a computer or other device based on recognition of hand gestures, there may not be value in accurately reconstructing the motion of any other type of object (e.g., if a cat walks through the field of view, it may be sufficient to determine that the moving object is not a hand).

15 [00149] Analysis techniques in accordance with embodiments of the present invention can be implemented as algorithms in any suitable computer language and executed on programmable processors. Alternatively, some or all of the algorithms can be implemented in fixed-function logic circuits, and such circuits can be designed and fabricated using conventional or other tools.

[00150] Computer programs incorporating various features of the present invention may be
20 encoded on various computer readable storage media; suitable media include magnetic disk or tape, optical storage media such as compact disk (CD) or DVD (digital versatile disk), flash memory, and any other non-transitory medium capable of holding data in a computer-readable form. Computer readable storage media encoded with the program code may be packaged with a compatible device or provided separately from other devices. In addition program code may be
25 encoded and transmitted via wired optical, and/or wireless networks conforming to a variety of protocols, including the Internet, thereby allowing distribution, e.g., via Internet download.

[00151] The motion capture methods and systems described herein can be used in a variety of applications. For example, the motion of a hand can be captured and used to control a computer system or video game console or other equipment based on recognizing gestures made by the hand.
30 Full-body motion can be captured and used for similar purposes. In such embodiments, the analysis and reconstruction advantageously occurs in approximately real-time (e.g., times comparable to human reaction times), so that the user experiences a natural interaction with the equipment. In other applications, motion capture can be used for digital rendering that is not done in real time, e.g., for computer-animated movies or the like; in such cases, the analysis can take as long as
35 desired. In intermediate cases, detected object shapes and motions can be mapped to a physical model whose complexity is suited to the application — i.e., which provides a desired processing speed given available computational resources. For example, the model may represent generic

5 hands at a computationally tractable level of detail, or may incorporate the user's own hands by initial image capture thereof followed by texture mapping onto a generic hand model. The physical model is manipulated ("morphed") according to the detected object orientation and motion.

[00152] Thus, although the invention has been described with respect to specific embodiments, it will be appreciated that the invention is intended to cover all modifications and equivalents within
10 the scope of the following claims.

[00153] In various embodiments, the system and method for capturing 3D motion of an object as described herein may be integrated with other applications, such as a head-mounted device or a mobile device. Referring to FIG. 29A, a head-mounted device 2902 typically includes an optical assembly that displays a surrounding environment or a virtual environment to the user;

15 incorporation of the motion-capture system 2904 in the head-mounted device 2902 allows the user to interactively control the displayed environment. For example, the virtual environment may include virtual objects that can be manipulated by the user's hand gestures, which are tracked by the motion-capture system 2904. In one embodiment, the motion-capture system 2904 integrated with the head-mounted device 2902 detects a position and shape of user's hand and projects it on
20 the display of the head-mounted device 2902 such that the user can see her gestures and interactively control the objects in the virtual environment. This may be applied in, for example, gaming or internet browsing.

[00154] Referring to FIG. 29B, in some embodiments, the motion-capture system 2904 is employed in a mobile device 2906 that communicates with other devices 2910. For example, a
25 television (TV) 2910 may include an input that connects to a receiver (e.g., a wireless receiver, a cable network or an antenna) to enable communication with the mobile device 2906. The mobile device 2906 first uses the embedded motion-capture system 2904 to detect movement of the user's hands, and to remotely control the TV 2910 based on the detected hand movement. For example, the user may perform a sliding hand gesture, in response to which the mobile device 2906 transmits
30 a signal to the TV 2910; the signal may be a raw trajectory that circuitry associated with the TV interprets, or the mobile device 2906 may include programming that interprets the gesture and sends a signal (e.g., a code corresponding to "sliding hand") to the TV 2910. Either way, the TV 2910 responds by activating and displaying a control panel on the TV screen, and the user makes selections thereon using further gestures. The user may, for example, move his hand in an "up" or
35 "down" direction, which the motion-capture system 2904 embedded in the mobile device 2906 converts to a signal that is transmitted to the TV 2910, and in response, the user's selection of a channel of interest from the control panel is accepted. Additionally, the TV 2910 may connect to a

5 source of video games (e.g., video game console or web-based video game). The mobile device 2906 may capture the user's hand motion and transmit it to the TV for display thereon such that the user can remotely interact with the virtual objects in the video game.

[00155] Referring to FIG. 29C, in various embodiments, the motion-capture system 2904 is integrated with a security system 2912. The security system 2912 may utilize the detected hand
10 shape as well as hand jitter (detected as motion) in order to authenticate the user 2914. For example, an authentication server 2916 may maintain a database of users and corresponding hand shapes and jitter patterns. When a user 2914 seeks access to a secure resource 2912, the motion-capture system 2904 integrated with the resource 2912 (e.g., a computer) detects the user's hand shape and jitter pattern and then identifies the user 2914 by transmitting this data to the
15 authentication server 2916, which compares the detected data with the database record corresponding to the access-seeking user 2914. If the user 2914 is authorized to access the secure resource 2912, the server 2916 transmits an acknowledgment to the resource 2912, which thereupon grants access. It should be stressed that the user 2914 may be authenticated to the secure system 2912 based on the shape of any part of a human body that may be detected and recognized
20 using the motion-capture system 2904.

[00156] The terms and expressions employed herein are used as terms and expressions of description and not of limitation, and there is no intention, in the use of such terms and expressions, of excluding any equivalents of the features shown and described or portions thereof. In addition, having described certain embodiments of the invention, it will be apparent to those of ordinary skill
25 in the art that other embodiments incorporating the concepts disclosed herein may be used without departing from the spirit and scope of the invention. Accordingly, the described embodiments are to be considered in all respects as only illustrative and not restrictive.

[00157] What is claimed is:

CLAIMS

- 1 1. A method of determining position and shape of an object in three-dimensional (3-D) space,
2 the method comprising:
3 obtaining one or more images of an object;
4 analyzing, by a computer, the one or more images to define at least four points on a surface
5 of the object in each one of a plurality of slices;
6 generating, by the computer, a cross-section of the object in each slice based on the at least
7 four points;
8 defining a 3-D model of the object based on the cross-sections in the plurality of slices; and
9 based on the 3-D model, determining, by the computer, a position and shape of the object.
- 1 2. The method of claim 1 wherein analyzing the one or more images to define the at least four
2 points includes, for at least one of the slices, defining at least four coplanar tangent lines to the
3 object in the slice.
- 1 3. The method of claim 1 wherein obtaining the one or more images of the object includes
2 using a time-of-flight camera to capture an image of the object and wherein analyzing the one or
3 more images to define the at least four points includes, for at least one of the slices, determining the
4 positions of the at least four points based on time-of-flight data provided by the time-of-flight
5 camera.
- 1 4. The method of claim 1 wherein defining the 3-D model of the object includes correlating
2 the cross-sections generated for each of the slices.
- 1 5. A method of determining position and shape of an object in three-dimensional (3-D) space,
2 the method comprising:
3 obtaining one or more silhouette images of an object;
4 analyzing, by a computer, the one or more silhouette images to define at least four coplanar
5 tangent lines to the object in each one of a plurality of slices;
6 generating, by the computer, a cross-section of the object in each slice based on the at least
7 four tangents;
8 defining a 3-D model of the object based on the cross-sections in the plurality of slices; and

9 based on the 3-D model, determining, by the computer, a position and shape of the object.

1 6. The method of claim 5 wherein obtaining the one or more silhouette images of the object
2 includes:

3 using at least two cameras, collecting at least two images of the object.

1 7. The method of claim 5 wherein obtaining the one or more silhouette images of the object
2 includes:

3 directing light from a light source toward the object; and

4 using at least one camera, collecting an image of the object and a shadow cast by the object.

1 8. The method of claim 5 wherein generating the cross-section includes generating the cross-
2 section as a simple closed curve.

1 9. The method of claim 5 wherein generating the cross-section includes generating the cross-
2 section as an elliptical cross-section.

1 10. The method of claim 9 wherein generating the cross-section includes, for at least one of the
2 slices:

3 initializing one parameter of an equation defining an ellipse to an assumed value; and

4 using the tangent lines and the initialized parameter, computing one or more complete
5 solution sets of parameters for the equation defining the ellipse.

1 11. The method of claim 10 wherein generating the cross-section further includes:

2 discarding any one of the one or more complete solution sets of parameters that does not
3 satisfy a physical constraint.

1 12. The method of claim 5 wherein defining the 3-D model of the object includes correlating
2 the cross-sections generated for each of the slices.

1 13. The method of claim 12 wherein defining the 3-D model includes:

2 determining an object type from the 3-D model; and

3 refining the cross-sections based on the object type.

1 14. A method for motion capture, the method comprising:
2 obtaining one or more silhouette images of a moving object at each of a plurality of times;
3 for at least one of the plurality of times, analyzing, by a computer, the one or more
4 silhouette images to define at least four coplanar tangent lines to the object in each one of a
5 plurality of slices;
6 generating, by the computer, a cross-section of the object in each slice based on the at least
7 four tangents;
8 constructing a 3-D model of the object based on the cross-sections in the plurality of slices;
9 based on the 3-D model, determining, by the computer, a position and a shape of the object
10 at the given time; and
11 repeating the acts of analyzing, generating and constructing for each of the plurality of times
12 to construct a model of a motion of the object.

1 15. The method of claim 14 further comprising:
2 correlating the determined position and shape of the object across different ones of the
3 plurality of times; and
4 refining the model of the motion of the object based on the correlation.

1 16. The method of claim 15 wherein refining the model of the motion of the object based on the
2 correlation includes eliminating from the model at a first time a cross-section that does not correlate
3 with the model at a second time.

1 17. The method of claim 14 further comprising:
2 determining, based on the 3-D model as constructed from images at a first one of the
3 plurality of times, an object type for the 3-D model; and
4 using the determined object type to constrain the construction of the 3-D model at a second
5 one of the plurality of times.

1 18. The method of claim 14 wherein the object includes two or more separately articulating
2 members and the model of the motion of the object includes a model of the motion of each of the
3 two separately articulating members.

1 19. A motion capture system comprising:

2 a camera subsystem; and
3 a processor coupled to receive image data from the camera subsystem, the processor being
4 configured to:
5 determine one or more silhouettes of an object from the image data; analyze the one
6 or more silhouettes to define at least four coplanar tangent lines to the object in each one of a
7 plurality of slices;
8 generate a cross-section of the object in each slice based on the at least four
9 tangents;
10 define a 3-D model of the object based on the cross-sections in the plurality of
11 slices; and
12 determine, based on the 3-D model, a position and shape of the object.

1 20. The motion capture system of claim 19 wherein the camera subsystem includes a first
2 camera and a second camera arranged at known positions and having overlapping fields of view.

1 21. The motion capture system of claim 19 wherein the camera subsystem includes:
2 a camera; and
3 a light source at a known position and configured to cast a shadow of an object into a field
4 of view of the camera,
5 wherein the camera is configured to obtain an image that includes both the object and the
6 shadow of the object.

1 22. The motion capture system of claim 21 wherein the processor is further configured to
2 determine the one or more silhouettes of the object by locating the object and the shadow of the
3 object in a single image obtained by the camera.

1 23. The motion capture system of claim 19 wherein the camera subsystem includes:
2 a camera; and
3 a plurality of light sources, each light source having a known position and being configured
4 to cast a shadow of an object into a field of view of the camera,
5 wherein the camera is configured to obtain an image that includes the shadows of the object
6 cast by the plurality of light sources.

- 1 24. The motion capture system of claim 19 wherein the camera subsystem includes at least one
2 infrared camera.
- 1 25. The motion capture system of claim 19 wherein the camera subsystem includes:
2 a camera;
3 a front-surface mirror; and
4 a beamsplitter disposed at an angle to the front-surface mirror,
5 wherein the first camera is oriented toward the beamsplitter and receives multiple images of
6 the object simultaneously, wherein the multiple images are created by light passing through the
7 beamsplitter and the front-surface mirror.
- 1 26. A method of identifying a position and shape of an object in three-dimensional (3D) space,
2 the method comprising:
3 capturing an image generated by casting an output from at least one source onto the object;
4 analyzing the image to computationally slice the object into a plurality of two-dimensional
5 (2D) slices, each slice corresponding to a cross-section of the object;
6 identifying shapes and positions of a plurality of cross-sections of the object based at least
7 in part on the image and a location of the at least one source; and
8 reconstructing the position and shape of the object in 3D space based at least in part on a
9 plurality of the identified cross-sectional shapes and positions of the object.
- 1 27. The method of claim 26, wherein the at least one source is at least one light source.
1
- 2 28. The method of claim 26, wherein the at least one source is at least one sonic source.
- 1 29. The method of claim 26, wherein the cross-sectional shape and position of the object is
2 identified by selecting a collection of intersection points generated by connecting the at least one
3 source to points on the image associated with the 2D slice.
- 1 30. The method of claim 29, wherein the intersection points are selected based on a total
2 number of the at least one source.
- 1 31. The method of claim 29, wherein the image is a shadow of the object.

- 1 32. The method of claim 31, wherein the intersection points are selected based on locations of
2 projection points associated with the intersection points, the projection points being projected from
3 the intersection points to the 2D slice using the at least one source.
- 1 33. The method of claim 31, further comprising:
2 splitting the cross-section of the object into multiple small regions and using each small
3 region to generate at least a portion of the shadow image of the 2D slice; and
4 identifying the small regions based on the shadow image of the 2D slice and the location of
5 the at least one source.
- 1 34. The method of claim 33, wherein the small regions are identified if a number of the
2 intersection points is equal to or greater than a predetermined criteria number.
- 1 35. The method of claim 29, wherein the intersection points are selected based on the location
2 of the at least one source and a size of the cross-section of the image.
- 1 36. The method of claim 29, wherein the selected collection of intersection points in a first 2D
2 slice is reused in a second 2D slice.
- 1 37. The method of claim 29, wherein the image includes reflections from the object.
- 1 38. The method of claim 37, wherein the intersection points are selected based on time-of-flight
2 data using a time-of-flight camera.
- 1 39. The method of claim 26, further comprising defining a 3D model of the object and
2 reconstructing the position and shape of the object in 3D space based on the 3D model.
- 1 40. The method of claim 26, wherein the position and shape of the object in 3D space is
2 reconstructed based on correlations between the plurality of the 2D slices.
- 1 41. The method of claim 26, further comprising defining coordinates of the image.

- 1 42. The method of claim 26, further comprising separating the image into a plurality of primary
2 images, each comprising a color.
- 1 43. The method of claim 42, further comprising identifying various colors on the object based
2 on the plurality of the primary images.
- 1 44. The method of claim 26, wherein the image is generated by casting light from a plurality of
2 light sources onto the object, the light sources being aligned in a line or in a plane.
- 1 45. The method of claim 26, wherein the object is a human or a human body part.
- 1 46. The method of claim 26, wherein the object is a handheld object such as a pencil or a
2 scalpel.
- 1 47. The method of claim 26, further comprising manipulating at least one virtual object
2 displayed on a device based on the identified position and shape of the object.
- 1 48. The method of claim 47, wherein the device is a head-mounted device.
- 1 49. The method of claim 47, wherein the device is a TV.
- 1 50. The method of claim 47, wherein the identified position and shape of the object is used to
2 manipulate the virtual object via wireless cell phone communication.
- 1 51. The method of claim 26, further comprising authenticating a user based on at least one of a
2 detected shape of the object or detected motion of the object and subsequent matching thereof to
3 data in a database record corresponding to the user.
- 1 52. A system for identifying a position and shape of an object in three-dimensional (3D) space,
2 the system comprising:
3 at least one camera oriented toward a field of view;
4 at least one source to direct illumination onto the object in the field of view; and
5 an image analyzer coupled to the camera and the at least one source and configured to:

6 operate the camera to capture at least one image of the object; and
7 identify a position and shape of the object in 3D space based on the captured
8 image and a location of the at least one source.

1 53. The system of claim 52, wherein the at least one source is at least one light source.

1 54. The system of claim 53, further comprising a plurality of filters placed on light sensors of
2 the camera to generate multiple images, each image corresponding to a color filter.

1 55. The system of claim 53, wherein the at least one light source is a plurality of light sources
2 each aligned in a line or in a plane.

1 56. The system of claim 52, wherein the at least one source is at least one sonic source.

1 57. The system of claim 52, wherein the image analyzer is further configured to:
2 slice the object into a plurality of two-dimensional (2D) slices, each slice corresponding to a
3 cross-section of the object;
4 identify a shape and position of the object based at least in part on an image captured by the
5 camera and a location of the at least one light source; and
6 reconstruct the position and shape of the object in 3D space based at least in part on a
7 plurality of the identified cross-sectional shapes and positions of the object.

1 58. The system of claim 52, wherein the camera is a time-of-flight camera.

1 59. The system of claim 52, wherein the image analyzer is further configured to define a 3D
2 model of the object and reconstruct the position and shape of the object in 3D space based on the
3 3D model.

1 60. The system of claim 52, further comprising a secondary device operatively connected to the
2 system.

1 61. The system of claim 60, wherein the secondary device is a head-mounted device.

1 62. The system of claim 60, wherein the secondary device is a mobile device.

1 63. The system of claim 60, wherein the secondary device is an authentication server for
2 authenticating a user based on at least one of a shape or a jitter of the user's hand detected by the
3 image analyzer.

1 64. A method of identifying a position and shape of an object in three-dimensional (3D) space,
2 the method comprising:
3 capturing, with a single camera, an image generated by casting an output from at least one
4 source onto the object;
5 analyzing the image to computationally slice the object into a plurality of two-dimensional
6 (2D) slices, each slice corresponding to a cross-section of the object, based at least in part on a
7 plurality of edge points in the image; and
8 reconstructing the position and shape of at least a portion of the object in 3D space based at
9 least in part on a plurality of the identified cross-sectional positions and sizes.

1 65. The method of claim 64, wherein the at least one source is a light source.

1 66. The method of claim 65, wherein the edge points define a viewed portion of the cross-
2 section within which the portion of the cross-section is within a field of view of an image-capturing
3 device.

1 67. The method of claim 66, wherein the at least one source is three light-emitting diodes.

1 68. The method of claim 66, wherein the at least one source is more than three light-emitting
2 diodes.

1 69. The method of claim 66, wherein the at least one source is two light-emitting diodes.

1 70. The method of claim 66, wherein light rays cast from the edge points to the image-capturing
2 device are tangent to the cross-section.

- 1 71. The method of claim 64, wherein at least one of the edge points is created by emitting light
2 from the at least one source onto the object.
- 1 72. The method of claim 71, wherein a light ray cast from the at least one source onto the at
2 least one of the edge points is tangent to the cross-section.
- 1 73. The method of claim 71, wherein the at least one of the edge points is defined by a
2 boundary between a shadow region and an illuminated region on the cross-section of the object.
- 1 74. The method of claim 64, further comprising defining a 3D model of the object and
2 reconstructing the position and shape of the object in 3D space based on the 3D model.
- 1 75. The method of claim 64, wherein the position and shape of the object in 3D space is
2 reconstructed based on correlations between the plurality of the 2D slices.
- 1 76. The method of claim 64, wherein the light source is a plurality of light sources and further
2 comprising operating the plurality of light sources in a pulsed fashion, whereby a plurality of the
3 edge points are generated sequentially.
- 1 77. A system for identifying a position and shape of an object in three-dimensional (3D) space,
2 the system comprising:
3 a single camera oriented toward a field of view;
4 at least one source to direct illumination onto the object in the field of view; and
5 an image analyzer coupled to the camera and the at least one source and configured to:
6 capture an image generated by casting an output from at least one source onto the
7 object;
8 analyze the image to computationally slice the object into a plurality of two-
9 dimensional (2D) slices, each slice corresponding to a cross-section of the object, based at
10 least in part on a plurality of edge points in the image; and
11 reconstruct the position and shape of at least a portion of the object in 3D space
12 based at least in part on a plurality of the identified cross-sectional positions and sizes.

1 78. The system of claim 77, wherein the at least one source is a plurality of light sources.

1 79. The system of claim 78, wherein the at least one source is three light-emitting diodes.

1 80. The system of claim 78, wherein the at least one source is more than three light-emitting
2 diodes

1 81. The system of claim 78, wherein the at least one source is two light-emitting diodes.

1 82. The system of claim 78, further comprising a driver for operating the plurality of light
2 sources in a pulsed fashion, whereby a plurality of the edge points are generated sequentially. The
3 system of claim 77, wherein the image analyzer is further configured to define a 3D model of the
4 object and reconstruct the position and shape of the object in 3D space based on the 3D model.




FIG. 1




FIG. 2




FIG. 3A




FIG. 3B




FIG. 4A




FIG. 4B




FIG. 4C




FIG. 5




FIG. 6A




FIG. 6B




FIG. 7

$$C_1x^2 + C_2xy + C_3y^2 - (2C_1x_c + C_2y_c)x - (2C_3y_c + C_2x_c)y + (C_1x_c^2 + C_2x_cy_c + C_3y_c^2 - 1) = 0$$

$$C_1 = \frac{\cos^2 \theta}{a^2} + \frac{\sin^2 \theta}{b^2}$$

$$C_2 = -2\cos\theta\sin\theta\left(\frac{1}{a^2} - \frac{1}{b^2}\right)$$

$$C_3 = \frac{\sin^2 \theta}{a^2} + \frac{\cos^2 \theta}{b^2}$$

FIG. 8




FIG. 9

$$A_1x+B_1y+D_1=0$$

$$A_2x+B_2y+D_2=0$$

$$A_3x+B_3y+D_3=0$$

$$A_4x+B_4y+D_4=0$$

FIG. 10A

$$r_{13} = \begin{bmatrix} A_1 & B_1 \\ A_3 & B_3 \end{bmatrix} \setminus \begin{bmatrix} -D_1 \\ -D_3 \end{bmatrix}$$

$$r_{23} = \begin{bmatrix} A_2 & B_2 \\ A_3 & B_3 \end{bmatrix} \setminus \begin{bmatrix} -D_2 \\ -D_3 \end{bmatrix}$$

$$r_{14} = \begin{bmatrix} A_1 & B_1 \\ A_4 & B_4 \end{bmatrix} \setminus \begin{bmatrix} -D_1 \\ -D_4 \end{bmatrix}$$

$$r_{24} = \begin{bmatrix} A_2 & B_2 \\ A_4 & B_4 \end{bmatrix} \setminus \begin{bmatrix} -D_2 \\ -D_4 \end{bmatrix}$$

FIG. 10B

$$c1=(r_{13}+r_{24})/2$$

$$c2=(r_{14}+r_{23})/2$$

$$\delta1=c2_1-c1_1$$

$$\delta2=c2_2-c1_2$$

$$p=\delta1/\delta2$$

$$q=c1_1-c1_2*p$$

$$G=A*p+B$$

$$H=A*q+D$$

FIG. 10C

$$v = \begin{bmatrix} G_2^2 & G_3^2 & G_4^2 \\ (G_2H_2)^2 & (G_3H_3)^2 & (G_4H_4)^2 \\ H_2^2 & H_3^2 & H_4^2 \end{bmatrix} \setminus \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

$$w = \begin{bmatrix} G_2^2 & G_3^2 & G_4^2 \\ (G_2H_2)^2 & (G_3H_3)^2 & (G_4H_4)^2 \\ H_2^2 & H_3^2 & H_4^2 \end{bmatrix} \setminus \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$$

$$v_{A2} = (v_1A_1)^2 + (v_2A_2)^2 + (v_3A_3)^2$$

$$v_{AB} = (v_1A_1B_1)^2 + (v_2A_2B_2)^2 + (v_3A_3B_3)^2$$

$$v_{B2} = (v_1B_1)^2 + (v_2B_2)^2 + (v_3B_3)^2$$

$$w_{A2} = (w_1A_1)^2 + (w_2A_2)^2 + (w_3A_3)^2$$

$$w_{AB} = (w_1A_1B_1)^2 + (w_2A_2B_2)^2 + (w_3A_3B_3)^2$$

$$w_{B2} = (w_1B_1)^2 + (w_2B_2)^2 + (w_3B_3)^2$$

FIG. 10D

$$_8t^6 + Q_7t^7 + Q_6t^6 + Q_5t^5 + Q_4t^4 + Q_3t^3 + Q_2t^2 + Q_1t + Q_0$$

FIG. 10E

$$4A_1^2n^2v_{B2}^2 + 4v_{B2}B_1^2(1-n^2v_{A2}) - (G_1(1-n^2v_{A2})w_{B2} + n^2v_{B2}w_{A2} + 2H_1v_{B2})^2$$

FIG. 10F

$$- (2(2n^2v_{AB}w_{A2} + 4H_1v_{AB} + 2G_1n^2v_{AB}w_{B2} + 2G_1(1-n^2v_{A2})w_{AB})) (G_1(1-n^2v_{A2})w_{B2} + n^2v_{B2}w_{A2} + 2H_1v_{B2}) - 8A_1B_1n^2v_{B2}^2 + 16A_1^2n^2v_{AB}v_{B2} + (4(2A_1B_1(1-n^2v_{A2}) + 2B_1^2n^2v_{AB}))v_{B2} + 8B_1^2(1-n^2v_{A2})v_{AB}$$

FIG. 10G

$$- (2(2H_1v_{B2} + 2H_1v_{A2} + n^2v_{A2}w_{A2} + n^2v_{B2}(-2w_{AB} + w_{B2}) + G_1(n^2v_{B2} + 1)w_{B2} + 4G_1n^2v_{AB}w_{AB} + G_1(1-n^2v_{A2})v_{A2})) \times (G_1(1-n^2v_{A2})w_{B2} + n^2v_{B2}w_{A2} + 2H_1v_{B2}) - (2n^2v_{AB}w_{A2} + 4H_1v_{AB} + 2G_1n^2v_{AB}w_{B2} + 2G_1(1-n^2v_{A2})w_{AB})^2 + 4B_1^2n^2v_{B2}^2 - 32A_1B_1n^2v_{AB}v_{B2} + 4A_1^2n^2(2v_{A2}v_{B2} + 4v_{AB}^2) + 4A_1^2n^2v_{B2}^2 + (4(A_1^2(1-n^2v_{A2}) + 4A_1B_1n^2v_{AB} + B_1^2(-n^2v_{B2} + 1) + B_1^2(1-n^2v_{A2})))v_{B2} + (8(2A_1B_1(1-n^2v_{A2}) + 2B_1^2n^2v_{AB}))v_{AB} + 4B_1^2(1-n^2v_{A2})v_{A2}$$

FIG. 10H

$$\begin{aligned}
& : - (2(4H_1v_{AB} + 2G_1(-n^2v_{B2}+1)w_{AB}+2G_1n^2v_{AB}v_{A2}+2n^2v_A(-2w_{AB}+w_{B2}))) (G_1(1-n^2v_{A2})w_{B2} + n^2v_{B2}w_{A2}+2H_1v_{B2})- \\
& (2(2H_1v_{B2}+2H_1v_{A2}+ n^2v_{A2}w_{A2} + n^2v_{B2}(-2w_{AB}+w_{B2}))+G_1(-n^2v_{B2}+1)w_{B2}+ 4G_1n^2v_{AB}w_{AB}+G_1(1-n^2v_{A2})v_{A2}))x \\
& (2n^2v_{AB}w_{A2}+4H_1v_{AB}+2G_1n^2v_{AB}w_{B2} + 2G_1(1-n^2v_{A2})w_{AB})+ 16B_1^2n^2v_{AB}v_{B2} -8A_1B_1n^2(2v_{A2}v_{B2}+4v_{AB}^2)+ \\
& 16A_1^2n^2v_{A2}v_{AB} - 8A_1B_1n^2v_{B2}^2 + 16A_1^2n^2v_{AB}v_{B2}+(4(2A_1^2n^2v_{AB}+ 2A_1B_1(- n^2v_{B2} +1)+ 2A_1B_1(1-n^2v_{A2}))+2B_1^2n^2v_{AB}))v_{B2} + \\
& (8(A_1^2(1-n^2v_{A2})+ 4A_1B_1n^2v_{AB}+B_1^2(-n^2v_{B2}+1)+B_1^2(1- n^2v_{A2})))v_{AB} + (4(2A_1B_1(1- n^2v_{A2}))+2B_1^2n^2v_{AB}))v_{A2}
\end{aligned}$$

FIG. 10I

$$\begin{aligned}
& = (4(A_1^2(-n^2v_{B2})+ A_1^2(1- n^2v_{A2}))+4A_1B_1n^2v_{AB}+B_1^2(- n^2v_{B2}+1)))v_{B2} + \\
& (8(2A_1^2n^2v_{AB}+ 2A_1B_1(-n^2v_{B2}+1)+2A_1B_1(1-n^2v_{A2}))+2B_1^2n^2v_{AB}))v_{AB} + \\
& (4(A_1^2(1-n^2v_{A2})+ 4A_1B_1n^2v_{AB}+B_1^2(-n^2v_{B2}+1)+B_1^2(1- n^2v_{A2})))v_{A2}+ 4B_1^2n^2(2v_{A2}v_{B2}+4v_{AB}^2)- 32A_1B_1n^2v_{A2}v_{AB}+ 4A_1^2n^2v_{A2}^2 + \\
& 4B_1^2n^2v_{B2}^2 - 32A_1B_1n^2v_{AB}v_{B2}+ 4A_1^2n^2(2v_{A2}v_{B2}+4v_{AB}^2) - \\
& (2(G_1(-n^2v_{B2}+ 1)v_{A2} + n^2v_{A2}(-2w_{AB} +w_{B2}))+2H_1v_{A2}))(G_1(1-n^2v_{A2})w_{B2} + n^2v_{B2}w_{A2}+2H_1v_{B2})- \\
& (2(4H_1v_{AB} + 2G_1(-n^2v_{B2}+1)w_{AB}+2G_1n^2v_{AB}v_{A2}+2n^2v_{AB}(-2w_{AB}+w_{B2})))x \\
& (2n^2v_{AB}w_{A2}+4H_1v_{AB}+2G_1n^2v_{AB}w_{B2} + 2G_1(1-n^2v_{A2})w_{AB})- \\
& (2H_1v_{B2}+2H_1v_{A2}+ n^2v_{A2}w_{A2} + n^2v_{B2}(-2w_{AB}+w_{B2}))+G_1(-n^2v_{B2}+1)w_{B2}+ 4G_1n^2v_{AB}w_{AB}+G_1(1-n^2v_{A2})v_{A2})^2
\end{aligned}$$

FIG. 10J

$$\begin{aligned}
& G_1(-n^2v_{B2}+1)v_{A2}+n^2v_{A2}(-2w_{AB}+w_{B2})+2H_1v_{A2})(2n^2v_{AB}w_{A2}+4H_1v_{AB}+2G_1n^2v_{AB}w_{B2}+2G_1(1-n^2v_{A2})w_{AB})- \\
& H_1v_{AB}+2G_1(-n^2v_{B2}+1)w_{AB}+2G_1n^2v_{AB}v_{A2}+2n^2v_{AB}(-2w_{AB}+w_{B2}))x \\
& v_{B2}+2H_1v_{A2}+n^2v_{A2}w_{A2}+n^2v_{B2}(-2w_{AB}+w_{B2})+G_1(-n^2v_{B2}+1)w_{B2}+4G_1n^2v_{AB}w_{AB}+G_1(1-n^2v_{A2})v_{A2}+ \\
& n^2v_{A2}v_{AB}-8A_1B_1n^2v_{A2}^2+16B_1^2n^2v_{AB}v_{B2}-8A_1B_1n^2(2v_{A2}v_{B2}+4v_{AB}^2)+16A_1^2n^2v_{A2}v_{AB}+(4(2A_1^2n^2v_{AB}+2A_1B_1(-n^2v_{B2}+1)))v_{B2}+ \\
& v_1(-n^2v_{B2}+1)+A_1(1-n^2v_{A2})+4A_1B_1n^2v_{AB}+B_1(-n^2v_{B2}+1))v_{AB}+ \\
& A_1^2n^2v_{AB}+2A_1B_1(-n^2v_{B2}+1)+2A_1B_1(1-n^2v_{A2})+2B_1^2n^2v_{AB}))v_{A2}
\end{aligned}$$

FIG. 10K

$$\begin{aligned}
& (-n^2v_{B2}+1)v_{B2}+(8(2A_1^2n^2v_{AB}+2A_1B_1(-n^2v_{B2}+1)))v_{AB}+ \\
& n^2(-n^2v_{B2}+1)+A_1^2(1-n^2v_{A2})+4A_1B_1n^2v_{AB}+B_1^2(-n^2v_{B2}+1))v_{A2}+4B_1^2n^2v_{A2}^2+4B_1^2n^2(2v_{A2}v_{B2}+4v_{AB}^2)- \\
& v_1B_1n^2v_{A2}v_{AB}+4A_1^2n^2v_{A2}^2-(2(G_1(-n^2v_{B2}+1)v_{A2}+n^2v_{A2}(-2w_{AB}+w_{B2})+2H_1v_{A2}))x \\
& v_{B2}+2H_1v_{A2}+n^2v_{A2}w_{A2}+n^2v_{B2}(-2w_{AB}+w_{B2})+G_1(-n^2v_{B2}+1)w_{B2}+4G_1n^2v_{AB}w_{AB}+G_1(1-n^2v_{A2})v_{A2}- \\
& v_{AB}+2G_1(-n^2v_{B2}+1)w_{AB}+2G_1n^2v_{AB}v_{A2}+2n^2v_{AB}(-2w_{AB}+w_{B2}))^2
\end{aligned}$$

FIG. 10L

$$\begin{aligned}
& (-n^2v_{B2}+1)v_{AB}+(4(2A_1^2n^2v_{AB}+2A_1B_1(-n^2v_{B2}+1)))v_{A2}+16B_1^2n^2v_{A2}v_{AB}-8A_1B_1n^2v_{A2}^2- \\
& n^2(-n^2v_{B2}+1)v_{A2}+n^2v_{A2}(-2w_{AB}+w_{B2})+2H_1v_{A2})(4H_1v_{AB}+2G_1(-n^2v_{B2}+1)w_{AB}+2G_1n^2v_{AB}v_{A2}+2n^2v_{AB}(-2w_{AB}+w_{B2}))
\end{aligned}$$

FIG. 10M

$$(-n^2v_{B2}+1)v_{A2}-(G_1(-n^2v_{B2}+1)v_{A2}+n^2v_{A2}(-2w_{AB}+w_{B2})+2H_1v_{A2})^2+4B_1^2n^2v_{A2}^2$$

FIG. 10N




FIG. 11A




FIG. 11B




FIG. 11C




FIG. 12




FIG. 13




FIG. 14




FIG. 15




FIG. 16A




FIG. 16B




FIG. 16C




FIG. 17




FIG. 18




FIG. 19




FIG. 20




FIG. 21




FIG. 22A





FIG. 22B




FIG. 22C




FIG. 23




FIG. 24

CONVERTING IMAGES OF SHADOWS TO SHADOW POINTS




FIG. 25A




FIG. 25B




FIG. 26A




FIG. 26B




FIG. 26C




FIG. 27




FIG. 28




FIG. 29A




FIG. 29B




FIG. 29C