




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification ⁶ : G06T 5/00</p>	<p>A1</p>	<p>(11) International Publication Number: WO 98/24058 (43) International Publication Date: 4 June 1998 (04.06.98)</p>
<p>(21) International Application Number: PCT/CA97/00907 (22) International Filing Date: 28 November 1997 (28.11.97) (30) Priority Data: 60/032,057 29 November 1996 (29.11.96) US (71) Applicant (for all designated States except US): LIFE IMAGING SYSTEMS INC. [CA/CA]; Suite 300, 195 Dufferin Avenue, London, Ontario N6A 1K7 (CA). (72) Inventors; and (75) Inventors/Applicants (for US only): FENSTER, Aaron [CA/CA]; 107 Ambleside Drive, London, Ontario N6G 4N9 (CA). DUNNE, Shane [CA/CA]; 401, 587 Talbot Street, London, Ontario N6A 2T2 (CA). (74) Agents: NASSIF, Omar, A. et al.; Gowling, Strathy & Henderson, Suite 4900, Commerce Court West, Toronto, Ontario M5L 1J3 (CA).</p>		<p>(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, HU, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).</p> <p>Published With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</p>

(54) Title: ENHANCED IMAGE PROCESSING FOR A THREE-DIMENSIONAL IMAGING SYSTEM



(57) Abstract

The present invention provides a depth-limited image processing means employed in combination with a three-dimensional ultrasonographic system whereby, from a model representing at least a portion of a three-dimensional volumetric array, a user selects a ray-depth from at least one face of the model thereby forming a "thick-plane" region of interest. All voxels within the "thick-plane" are processed successively according to a user selected imaging technique along the depth of the ray. The results of the processing are projected onto corresponding screen pixels to produce an enhanced three-dimensional image. In another embodiment, an adaptive depth-limited image processing means is employed whereby voxels are first checked against a user selected condition. Voxels meeting this condition are then processed as above to produce an enhanced three-dimensional image. The user selected imaging technique may include, but is not limited to, Texture Mapping, Maximum Intensity Projection (MIP), Conditional MIP (CMIP) and/or Ray Summation (Raysum).

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakistan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

ENHANCED IMAGE PROCESSING FOR A THREE-DIMENSIONAL
IMAGING SYSTEM

FIELD OF THE INVENTION

5

The present invention relates to medical diagnostic imaging and in particular to a method and system for enhancing three-dimensional images.

BACKGROUND OF THE INVENTION

10

In the medical field, it is common to use ultrasound diagnostic equipment to view internal organs of a subject. For example, in diagnosing Prostate cancer, a diagnostician uses transrectal ultrasound (TRUS) to identify whether lesions are present as well as to determine the location, size and extent of lesions if present. Conventional ultrasound diagnostic equipment typically comprise an ultrasound probe for transmitting ultrasound signals into the subject and receiving reflected ultrasound signals therefrom. The reflected ultrasound signals received by the ultrasound probe are processed and a two-dimensional image of the target under examination is formed.

20

Unfortunately, this conventional equipment produces two-dimensional images even though the target under examination is three-dimensional. Also, the two dimensional images represent a single thin plane taken at an arbitrary angle

to the target making it very difficult to localize the image plane in the target and very difficult to reproduce an image at a particular location at a later time. Furthermore, the imaged plane is usually not a clear detailed image but rather a series of shaded shapes which requires a considerable level of skill to interpret.

5

In U.S. application serial no. 08/419,049 and U.S. Patent No. 5,454,371, assigned to the assignee of the present application, the contents of which are incorporated herein by reference, three-dimensional ultrasound imaging systems are described. Although, these systems overcome disadvantages associated with
10 the prior art, improvements to enhance imaging and to increase the speed by which three-dimensional images can be generated from two-dimensional ultrasound images are continually being sought.

It is therefore an object of the present invention to provide a novel system
15 and method for generating a three-dimensional image from a succession of two-dimensional images, a novel image processing means for enhancing the three-dimensional image viewed by a user and, a novel ultrasound imaging system.

20

SUMMARY OF THE INVENTION

According to one aspect of the present invention there is provided a three-dimensional imaging system for acquiring a succession of two-dimensional images of a target volume represented by an array of pixels $I(x,y,z)$ into a three-dimensional image represented by a volumetric image array $V(x,y,z)$ comprising:

scanning means to scan said target volume along a predetermined geometric scanning path and generate a succession of digitized two-dimensional images thereof representing cross-sections of said target volume on a plurality of planes spaced along said scanning path;

memory storing said succession of digitized two-dimensional images together with other related image data defining the location of said two-dimensional images in said memory and defining interpretation information relating to the relative position of pixels within said two-dimensional images and to the relative position of pixels in adjacent two-dimensional images within said target volume;

transformation means receiving said digitized two-dimensional images and said other related image data and transforming said two-dimensional images and said other related image data into a volumetric image array;

display means to generate a three-dimensional image of said target volume from said volumetric image array; and,

an image processing means for processing said three-dimensional image of said target volume to render an enhanced three-dimensional image.

5

Preferably, the image processing means includes user changeable image processing parameters.

Also preferably, the image processing means includes user selectable
10 image processing parameters.

Also preferably the changeable image processing parameters include modifying a ray depth.

15 Also preferably, the selectable image processing parameters include selecting from a plurality of image processing techniques.

Also preferably the image processing techniques include, but are not limited to, Texture Mapping, Maximum Intensity Projection (MIP), Conditional
20 MIP (CMIP) and/or Ray Summation (Raysum).

Also preferably, the user selectable and user changeable image processing parameters are selected and changed via a user interface.

Also preferably, the user interface is a graphical user interface.

5

In one embodiment, a depth-limited image processing means is employed whereby, from a model representing at least a portion of the three-dimensional volumetric array, a user selects a ray-depth from at least one face of the model thereby forming a "thick-plane" region of interest. All voxels within the "thick-
10 plane" are processed successively according to a user selected imaging technique along the depth of the ray, the results of which are projected onto corresponding screen pixels to produce an enhanced three-dimensional image.

In another embodiment, an adaptive depth-limited image processing
15 means is employed whereby, from a model representing at least a portion of the three-dimensional volumetric array, a user selects a ray-depth from at least one face of the model. In combination with the ray depth, a user specified condition is selected with which voxels are examined. Voxels within the model interior are first processed according to the user specified condition. Voxels meeting the
20 user specified condition are then processed by a user selected imaging technique, the results of which are projected onto corresponding screen pixels to produce

an enhanced three-dimensional image. All other voxels are ignored thereby increasing processing speed.

The ray-depth thereby limits the processing depth within the volumetric
5 array and therefore limits the number of voxels processed to produce the
enhanced three-dimensional image. Processing the reduced number of voxels
results in quick processing speeds which permits at or near real-time enhanced
three-dimensional image updating. Accordingly, this further permits the user
to view enhanced three-dimensional images as the model is manipulated (rotated,
10 translated etc.) via a user input device and graphical user interface.

The user selected imaging technique may include, but is not limited to,
Texture Mapping, Maximum Intensity Projection (MIP), Conditional MIP
(CMIP) and/or Ray Summation (Raysum).

15

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the present invention will now be described more fully
with reference to the accompanying drawings in which:

20 Figure 1 is a perspective view of a three-dimensional ultrasound imaging
system;

Figure 2 is a perspective view of an ultrasound probe actuating assembly forming part of the system illustrated in Figure 1;

Figure 3 is a side view of the probe actuating assembly illustrated in Figure 2 supporting an ultrasound probe;

5 Figure 3a is a perspective view of the geometry of the ultrasound probe movement when driven by the ultrasound probe actuating assembly;

Figure 3b is an illustration of a number of tilted two-dimensional images;

Figure 3c is an illustration of a shear transformation;

Figure 4 is a block diagram showing various hardware and software
10 modules of the computer forming part of the system illustrated in Figure 1;

Figure 5a is a flowchart showing some of the operational steps of the system illustrated in Figure 1;

Figure 5b is a flowchart showing additional operational steps of the system illustrated in Figure 1; and

15 Figure 6 shows a flowchart of steps performed by the display module during initialization;

Figures 7a to 7d show flowcharts of steps performed by the user interface and display modules when manipulating a displayed three-dimensional image;

Figures 8a to 8c show a three-dimensional image and model undergoing
20 a rotation about a vertical axis;

Figures 9a to 9c show a three-dimensional image and model undergoing a rotation in a direction from top-left towards bottom right about an axis, angled at about 30° to the horizontal and sloping up and to the right;

Figures 10a to 10c show a three-dimensional image and model in which
5 a plane of the model is translated towards the geometric center of the model;

Figures 11a to 11c show a three-dimensional image and model in which a plane of the model is rotated about an axis, angled at about 30° to the horizontal and sloping up and to the right;

Figures 12a to 12d show a three-dimensional image and model in which
10 a plane of the model is translated away from the geometric center of the model until it disappears;

Figure 13 shows a typical full screen display including a main display window and a control display window;

Figure 14 shows another full screen display further including a measure
15 control window and an animation control window;

Figures 15a to 15c show full screen displays further including an orientation view window;

Figure 16 shows a three-dimensional image produced using of a volumetric image array, employing conventional Maximum Intensity Projection
20 (MIP) and having the model outlines suppressed from the display;

Figure 17 shows a three-dimensional image and model produced by texture mapping, from a volumetric image array (color Doppler ultrasound of the carotid artery of a healthy volunteer) containing a mix of colored voxels representing flowing blood and gray scale voxels representing stationary tissue;

5 Figure 18 shows a three-dimensional image produced using the volumetric image array of Figure 17, but using Conditional MIP to project only the color information. The front face of the model has been translated outward to reveal the entire blood vessel structure;

Figure 19 shows a three-dimensional image produced using the model
10 and volumetric image array of Figure 16, employing a conventional Raysum algorithm;

Figure 20a shows a three-dimensional image and model produced by texture mapping, from a volumetric image array using power Doppler ultrasound of blood vessels in the spleen of a healthy volunteer;

15 Figure 20b shows the image of Figure 20a after translating the front model plane inward by a few millimeters;

Figure 20c shows the image of Figure 20b after translating the front model plane further inward by a few millimetres;

Figure 21a shows a three-dimensional image produced employing the
20 model and volumetric image array of Figure 16, using Depth-Limited Maximum Intensity Projection (DL-MIP) with a depth of about two millimetres;

Figure 21b shows the image of Figure 21a, having the depth of the DL--MIP operation doubled to about four millimetres;

Figure 21c shows the image of Figure 21a, having the depth of the DL--MIP operation tripled to about six millimetres;

5 Figure 22a shows the image of Figure 21c, having the front face of the model translated inward by a few millimetres, revealing structures at greater depth;

Figure 22b shows the image of Figure 22a, having the front face of the model has translated further inward;

10 Figure 23 shows a three-dimensional image produced from the volumetric image array of Figures 20a through 22b, with the model adjusted to have two exterior faces through which the DL-MIP image can be seen;

Figure 24 shows the image of Figure 23, having the image produced using a Depth-Limited Raysum (DL-Raysum) algorithm;

15 Figure 25 shows the image of Figure 16, produced using a Adaptive Depth-Limited MIP (ADL-MIP) algorithm;

Figure 26 shows the image of Figure 24, produced using a Adaptive Depth-Limited Raysum (ADL-Raysum) algorithm;

20 Figure 27 shows the image of Figure 26, but with an additional shading step having been performed to enhance perception of depth in the three-dimensional image;

Figure 28 shows the image of Figure 25, but with an additional shading step having been performed to enhance perception of depth in the three-dimensional image; and,

Figure 29 shows a schematic representation of a raycasting technique in accordance with one embodiment of the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

In the applicant's co-pending U.S. Patent applications serial nos. 10 08/419,049 and 08/562,590 and applicant's U.S. patent issued no. 5,454,371, three-dimensional ultrasound imaging systems are disclosed. In these systems, when a succession of two-dimensional images have been captured and digitized the two-dimensional images are stored as a stack to form an image data array. Before a three-dimensional image of the scanned target volume can be created, 15 the image data array must be reconstructed to form a volumetric image array. The process of reconstructing the image data array to the volumetric image data array is a time consuming process which results in significant delays before acquired two-dimensional images can be used to generate a three-dimensional image. The present invention relates to a three-dimensional ultrasound imaging 20 system which overcomes this problem by allowing a three-dimensional image to

be generated using the two-dimensional image data array without requiring the image data array to undergo reconstruction.

Referring now to Figure 1, a three-dimensional ultrasound imaging system in accordance with the present invention is shown and is generally indicated by reference numeral 20. The system 20 is capable of generating a three-dimensional ultrasound image of a target volume of a subject under examination from a succession of two-dimensional ultrasound images of the target volume and allow the generated three-dimensional image to be manipulated. The subject under examination may be inanimate or animate. In the later case, the system 20 may be used in both medical and veterinary environments and may be used as a diagnostic tool or during surgery to provide updated images of the target volume of the subject undergoing surgery.

The system 20 includes an ultrasound probe actuating assembly 22 for removably retaining an ultrasound probe 24. In this embodiment, the probe actuating assembly 22 is designed to move the ultrasound probe along a linear scanning path Z so that the succession of two-dimensional images of the target volume can be taken.

20

The ultrasound probe 24 is connected to a clinical ultrasound machine 28 via a communication line 30. The ultrasound machine 28 in turn is connected to a computer 32 via communication line 34. The computer 32 includes a keyboard (not shown), a monitor 36 with a display screen 36a and a graphical input device 38 such as a single button mouse. It should however be realized that many other graphical input devices can be used to allow a user to input commands to the computer. The computer 32 provides output signals to a controller 40 via communication line 42 which in turn provides control signals to the probe actuating assembly 22 via communication line 44 to control the scanning motion of the ultrasound probe 24.

Referring now to Figures 2 and 3, the ultrasound probe 24 and the probe actuating assembly 22 are better illustrated. As can be seen, the probe actuating assembly 22 includes a housing 22a accommodating a motor 22b coupled to a threaded output shaft 22c by way of a pair of reduction wheels 22d and 22e and a belt 22f. An I-block 22g is mounted to the threaded shaft 22c via an internal threaded hole so that the I-block 22g moves along the linear scanning path Z in either direction as represented by arrow B when the threaded shaft 22c is rotated via the motor 22b. A probe holder 22h is secured to the I-block 22g by way of screws 22i and 22j and supports the ultrasound probe 24. The probe holder 22h is positioned relative to the threaded shaft 22c so that the longitudinal axis of the

ultrasound probe 24 forms an angle α with an axis normal to the longitudinal axis of the threaded shaft 22c. The angle of inclination of the probe holder 22h can be adjusted by loosening the screws 22i and 22j, rotating the probe holder 22h to the desired position and tightening the screws.

5

The probe actuating assembly 22 is particularly suited to applications where ultrasound images of a subject's internal organs or lesions such as breast tumours within the trunk of a subject P are to be taken. In these instances, a layer of coupling

10 gel 22k is typically placed on the subject P between the subject and the ultrasound probe 24.

During normal operation, the motor 22b is operated so that it rotates at a constant velocity. Thus, the threaded shaft 22c is rotated at a constant rate
15 causing the probe holder 22h to move the ultrasound probe 24 at a constant rate along the linear scanning path Z. As the ultrasound probe 24 moves along the linear scanning path Z, it transmits ultrasound signals at specific predefined intervals which impinge on the target volume. Reflected ultrasound signals from the target volume are also received by the probe 24 and are converted into
20 analog signals by a crystal (not shown) in the ultrasound probe 24. These analog signals are conveyed to the clinical ultrasound machine 28 where a succession

of two-dimensional analog images of the target volume are generated. Each two-dimensional image represents a cross-section of the target volume having x and y dimensions (see Figure 3a). The operation of the ultrasound probe 24 and clinical ultrasound machine 28 is well known to those of skill in the art and therefore, will not be described in any further detail herein.

Since the velocity of the ultrasound probe 24 moving along the linear scanning path Z and the ultrasound signal transmit intervals of the ultrasound probe 24 are known, the relative position of the two-dimensional analog images of the target volume along the z-axis can be readily determined.

The two-dimensional analog images generated by the ultrasound machine 28 are conveyed to the computer 32 via communication line 34. The computer 32 in turn digitizes the two-dimensional images and stores the digitized image data together with other related image data in a manner which allows a three-dimensional image of the target volume to be displayed virtually without delay as will be described. Once displayed, the computer 32 allows the image to be manipulated as will also be described.

Referring now to Figure 4, a block diagram of the computer 32 is shown illustrating some of the hardware and software modules therein. As can be seen,

the computer 32 includes a frame grabber module 80, such as for example, an IMAXX Video Capture Board manufactured by Precision Digital Images Corporation of Redmond, Washington, to process the two-dimensional analog images received from the clinical ultrasound machine 28 via communication line 5 34. Specifically, the frame grabber module 80 captures and digitizes the succession of two-dimensional analog images. Once digitized, the succession of two-dimensional images is stored in local physical memory 82.

The computer 32 also includes a user interface module 84 for interpreting 10 input commands received via the graphical input device 38. As should be realized, the user interface module 84 controls and co-ordinates the operation of the other modules of system 20 in response to input from the graphical input device 38 allowing the user to control the system as desired.

15 Once a succession of two-dimensional images of the target volume has been captured and digitized by frame grabber module 80 and stored in the physical memory 82, the digitized information can be processed in a number of ways depending on the input commands received by the user interface module 84 from the graphical input device 38. Specifically, the digitized information 20 can be transferred to an external file storage memory 88. The digitized information, whether stored in local physical memory 82 or in the external file

storage memory 88, may be processed by a display module 92 in response to input received from graphical input device 38 so that a three-dimensional image of the target volume can be displayed on the screen 36a of the monitor 36 and manipulated as will be described further herein.

5

The computer 32 also includes a probe scanning control module 98 which provides output signals to controller 40 to actuate the probe actuating assembly 22 and sweep the ultrasound probe 24 along the linear scanning path Z as desired. The probe scanning control module 98 also receives input from the user interface module 84.

Three-Dimensional Imaging

Because the ultrasound probe 24 is swept along a linear path Z and since the ultrasound probe 24 is at an angle α relative to an axis normal to the longitudinal axis of the linear scanning path Z, the ultrasound images taken by the ultrasound machine are cross-sections (or image slices) of the target volume on parallel planes tilted by the angle α (see Figure 3b). Thus, by nature of the acquisition of the ultrasound images, there is a dependency relationship between the y and z coordinates which mathematically can be expressed by a shear transformation (see Figure 3c). Shear transformations, which are well known

to those of skill in the art of computer graphics, are linear. The linearity of both the scanning path Z and the shear transformation allows the display module 92 to use the shear transformation as a component of the viewing transformation to convert pixel co-ordinates into corresponding voxel coordinates. As such, the display module 92 can act directly on the acquired digitized two-dimensional ultrasound image data and generate a three-dimensional image without requiring the acquired ultrasound image data to undergo the reconstruction processes described in Applicant's U.S. Patent No. 5,454,371 and patent application serial no. 08/419,049. This significantly increases the speed by which a three-dimensional image may be displayed.

In order for the display module 92 to be able to act directly on the acquired digitized ultrasound image data, the display module 92 requires other related image data which includes:

1. an address pointer defining the address of the location in the physical memory 82 at which the acquired ultrasound image data starts;
2. the extent of the acquired two-dimensional image data including the number of pixels along the x and y axis of each two-

dimensional image as well as the number of two-dimensional images taken in the z-direction;

3. the physical distance between the centers of adjacent pixels in both
5 the x and y directions in each two-dimensional image as well as the physical distance between corresponding pixels in adjacent two-dimensional images (the z-distance); and
4. the tilt angle α .

10

The other related image data referred to at points 1 and 2 above defines the location of the acquired digitized ultrasound image data in the memory. This information allows the display module 92 to locate the acquired ultrasound image data and associate sections of the image data with the proper image slices
15 of the target volume. The other related image data referred to at points 3 and 4 above defines image data interpretation information which allows the display module 92 to interpret the acquired ultrasound image data correctly.

Before acquiring two-dimensional images of a target volume using system
20 20, the other related image data must be defined. Specifically, the frame grabber module 80 is programmed to write the acquired digitized ultrasound

image data to the appropriate locations of the physical memory 82 and to generate the address pointer which in turn is stored in a calibration file in physical memory 82. In order to generate the other related image data at step 2, the system 20 is calibrated. Specifically, during calibration, the velocity of the ultrasound probe 24 along the linear scanning path Z and the ultrasound signal transmit interval are determined so that the number of two-dimensional image slices of the target volume to be taken are known. Also, the number of pixels along the x and y axis of the two-dimensional images are measured. The number of pixels in the x and y directions defines the edges of each two-dimensional ultrasound image. Thus, these numbers are used to locate the ultrasound image data within the scanned target volume. Once the numbers are determined they are also stored in the calibration file.

During the system calibration, in order to generate the other related image data at step 3, the center to center distance between two pixels in the same line of an ultrasound image (ie. in the x direction) are determined, the center to center distance between adjacent pixels in two different lines of the ultrasound image are determined (ie. in the y direction) and the center to center distance between corresponding pixels in two adjacent ultrasound images are determined. During the determination of the distances between the pixels mentioned above, only a few sample measurements between two pixels in the x-direction of an

ultrasound image, two pixels in the direction of the same ultrasound image are taken and average distances are determined. Since the velocity of the ultrasound probe 24 along the linear scanning path Z is constant and since the ultrasound signal transmit interval is known, the distance between corresponding pixels in adjacent two-dimensional ultrasound images is easily calculated. Once these distances are determined, the distance values are stored in the calibration file. The tilt angle α of the ultrasound probe 24 is also measured and is stored in the calibration file. Once the above other related image data is stored in the calibration file in physical memory 82, the calibration file is stored in the external file storage memory 88. At this stage, the system 20 is ready to acquire two-dimensional image slices of the target volume as will now be described.

Image Capturing

With reference now to Figures 5a and 5b, when it is desired to operate the three-dimensional ultrasound imaging system 20 to acquire two-dimensional images of the target volume, assuming the system 20 has been calibrated and has the other related image data described above stored in the external file storage memory 88, the ultrasound probe 24 and probe actuating assembly 22 must be properly located with respect to the subject so that the ultrasound waves

transmitted by the ultrasound probe 24 are directed at the target volume (see block 102).

Initially, the ultrasound probe transmits ultrasound signals to scan a cross-section of the target volume. The reflected ultrasound signals received from the target volume are conveyed to the clinical ultrasound machine 28 wherein a two-dimensional analog image slice of the target volume upon which the ultrasound signals impinged, is created. The two-dimensional analog image is then conveyed to the computer 32 via communication line 34 wherein it is captured and digitized via frame grabber module 80. The digitized two-dimensional image is then stored in the memory 82.

A copy of the digitized two-dimensional image is then conveyed to the user interface module 84 and the frame is drawn on the screen 36a of the monitor 36 (block 104). The user then manually moves the probe 24 along the linear scanning path Z while it is transmitting ultrasound signals so that two-dimensional analog images generated by the clinical ultrasound machine 28 are captured and digitized by the frame grabber module 80. These two-dimensional images are also then drawn on the screen 36a of monitor 36 via user interface module 84 (block 106). Next, the user is prompted to confirm that the

ultrasound signals are properly directed at the target volume after having viewed the frames drawn on the screen 36a of the monitor (block 108). If the target volume is outside of the drawn frames, then operation returns to block 104. Otherwise, the user provides input to the user interface module 84 using the graphical input device 38 to signify that the target volume is within the drawn frames.

Once this has been done, with the probe actuating assembly 22 in place (block 110), the system 20 is ready for proper operation. When a user inputs a start command such as for example by using an external foot or hand switch (not shown), by selecting an appropriate icon displayed on the screen 36a using the graphical input device 38 or alternatively by using a voice command, the user interface module 84 signals the probe scanning module 98. Within the context of the present invention, icon refers to any graphical element displayed on the screen 36a which can be selected using graphical input device 38.

At this point in time, the probe scanning module 98 conveys control signals to the probe actuating assembly 22 via controller 40 so that the ultrasound probe 24 is moved along the linear scanning path Z at a constant velocity. While this occurs, the ultrasound probe 24 is conditioned to transmit ultrasound signals at the predefined intervals so that the entire target volume is

scanned. As the ultrasound probe receives reflected ultrasound signals, it conveys analog information to the clinical ultrasound machine 28 which in and generates two-dimensional analog images. In this manner, a succession of two-dimensional analog images of the target volume representing a volume image are
5 generated by the clinical ultrasound machine 28 in response to the output of the ultrasound probe 24 (block 112). The succession of two-dimensional analog images generated by the clinical ultrasound machine 28 are captured and digitized by the frame grabber module 80. The digitized two-dimensional images are then conveyed to the physical memory 82 and stored as a stack to
10 form an array of two-dimensional images $I(x,y,z)$ with the pixels in the array $I(x,y,z)$ representing pixels of the digitized two-dimensional images. Because the computer 32 controls the movement of the probe actuating assembly 22 and the operation of the ultrasound probe 24, the spatial orientation of the individual two-dimensional images relative to the target volume is known.

15

In the present embodiment the two-dimensional images are considered to be grayscale images. However, the present invention does not depend on the "color" of the two-dimensional images to function properly. A grayscale pixel is associated with a gray-level having a value between 0 and $(2^n - 1)$ inclusively,
20 with n being the number of bits required for storing the gray-levels. The gray-level 0 is usually used as a "background color" and is said to be Black.

Once the two-dimensional images have been acquired and saved in physical memory 82 to form array $I(x,y,z)$, the user interface module 84 generates a prompt to signify that this stage of the image capturing has been completed. At this time, the user may review the acquired frames individually

5 in the manner described previously (block 114). If the two-dimensional images have been acquired incorrectly (block 116), the user can condition the system 20 to return to block 102. If the two dimensional images have been acquired correctly, the user interface module 84 generates a prompt to determine if the acquired two-dimensional images together with the other related image data in

10 the calibration file are to be saved in the external file storage memory 88 (block 120) together as a sheared volumetric image array $V(x,y,z)$. If the user selects the prompt, the acquired two-dimensional images together with the other related image data are saved as the sheared volumetric image array in the external file storage memory 88 (block 124).

15

Following this, or if the user does not elect to save the two-dimensional image data to the external memory 88, the user is prompted to decide whether the three-dimensional image is to be displayed on the screen 36a of the monitor 36 (block 126). If the user wishes to view the three-dimensional image, and the

20 image data is stored in the external file storage memory 88 in the form of a sheared volumetric image array, then the display module 92 retrieves the sheared

volumetric image array from the external file storage memory. If the image data has not been stored in the external file storage memory 88, the display module 92 retrieves the image data stored in the memory 82 and retrieves the calibration file from the external file storage memory 88. In either case, once the display
5 module 92 receives the image data and other related image data, it uses the data in conjunction with the shear transformation to display a three-dimensional image of the target volume on the screen 36a (block 128) with virtually no delay. The displayed image can be manipulated by the user as will be described. During image manipulation, the user can store displayed views in the memory
10 82 or in the external file storage memory 88 so that these views can be retrieved and re-examined at a later time. Once image manipulation has been completed, the user is prompted to confirm whether another three-dimensional image is to be created (block 130). If the user wishes to create another three-dimensional image, the system 20 reverts to block 102. Otherwise, the three-dimensional
15 imaging procedure is considered to be completed. If at block 126, the user does not elect to view the three-dimensional image, the system proceeds directly to block 130.

Three-Dimensional Image Display

20

Once the two-dimensional images of the target volume have been acquired and have been stored either in external file storage memory 88 with the calibration file as a sheared volumetric image array or in physical memory 82, the user can decide whether the three-dimensional image is to be displayed on the screen of monitor 36 using the graphical input device 38 (block 126). When the three-dimensional image is to be displayed as signified by the user via the graphical input device 38, the display module 92 enters an initialization routine (see Figure 6). Upon entering the initialization routine, the display module 92 sets a number of parameters to their default values as will be explained and uses this information when an image is to be displayed on the screen 36a (blocks 302 to 306). The parameters which are set to their default value are mouse-drag sensitivity, successive refinement sequence and pseudo-color mapping.

After the above has been done, the display module 92 determines whether the image data has been stored with the calibration file as a sheared volumetric image array or resides in the physical memory 82. If the image data has been stored as a sheared volumetric image array, the display module 92 retrieves a copy of the sheared volumetric image array from external file storage memory 88 (block 308). If the image data resides in the physical memory 82, the display module retrieves the image data from the memory 82 and retrieves the other related image data in the calibration file from the memory 88.

Once the display module 92 has retrieved the image data and the other related image data, the display module 92 checks to see if the image data has been assigned a model (block 310). The model is in the form of a convex polyhedron having a plurality of planar faces defined in the same space as the
5 three-dimensional image to be displayed. The polyhedron may take a variety of shapes as will be described herein.

If the image data has not been assigned a model, the dimensions of the image data are examined and a model in the form of a parallelepiped, which
10 encloses substantially all of the pixels in the image data array is created and assigned to the image data automatically (block 312). The parallelepiped has two interior angles at right angles and a third interior angle equal to the angle α . The assigned model is then saved in memory with the image data array. After this, the image data array is examined to determine if any or all preferred
15 Views A to C have been assigned to the image data array (block 314). If some or all of the preferred Views have not been assigned, the preferred Views are created automatically and saved in the memory 82 with the image array (block 315).

20 The model is then projected on the screen of the monitor 36 within a rectangular sub-region of the full screen display, henceforth called the "main

display window" via an orthographic projection (block 316). Figure 13 illustrates the screen 36a of monitor 36 on which the model and three-dimensional image are displayed " in the main display window MDW. It should have been realized that other projection transformations such as a perspective projection may be used to project the model onto the screen. only the visible faces of the model are displayed on the screen 36a, i.e., hidden-surface elimination is performed so that the displayed model has an opaque appearance. The perimeter lines of the model faces are depicted on the screen by white lines. Each visible model face is projected onto the screen within a polygonal area. Each visible face's polygonal area is divided into an internal area and an external area, the latter being that part of the total area immediately adjacent to the displayed boundary of the face and represents a small portion of the total area.

The display of each visible model face is accomplished as follows. Each screen pixel within the polygonal area of the displayed face has an associated two-dimensional Cartesian coordinate pair (x',y') . With precise specifications of the model surface, this can be extended to a three-dimensional coordinate triple (x',y',z') .

By means of the shear transformation, the pixel co-ordinates (x',y',z') may be converted to corresponding voxel coordinates (x,y,z) , to select a voxel

value within the volumetric image array $V(x,y,z)$. The extracted voxel value is indexed into the pseudo-color mapping table to yield a gray-level or color. The gray-level or color in turn is used to illuminate the screen pixel. This process is repeated for all screen pixels located within the displayed model faces (block 5 318) to render the image. This technique of display is called "texture mapping", and is known to those of skill in the art.

In addition to texture-mapping, the preferred embodiment of the present invention supports a variety of conventional three-dimensional volumetric image processing means, including Maximum Intensity Projection (MIP) and Ray summation hereinafter referred to as Raysum, (also typically known as Transparent Reprojection or Synthetic Radiograph to those skilled in the art). However, as will be described in greater detail below, the present inventors have developed several novel image processing means which are improvements to the MIP and Raysum rendering techniques. Specifically, such novel image processing means are Depth-Limited (DL) and Adaptive Depth-Limited (ADL) rendering which may be applied to both MIP and Raysum methods to enhance imaging. Furthermore, to enhance the user's perception of a three-dimensional structure in the resulting images, shading techniques including Depth Shading and Surface Shading may be applied to the results of the DL and ADL techniques. These methods will now be described.

The various rendering techniques previously mentioned can most easily be described in terms of the paradigm of "ray casting", known to those skilled in the art of three-dimensional image processing. In the preferred embodiment of the present invention, actual implementation of the DL and ADL techniques follow the ray casting paradigm, but it should be realized that other implementation techniques which provide substantially equivalent output may be substituted. For example, when fast texture mapping hardware is available, implementation based on front-to-back slice compositing as described by U Cullip and U Neumann, "Accelerating Volume Reconstruction with 3D Texture Hardware" UNC Tech Report, 1993" and P.A. Sabella, "Rendering Algorithm for Visualizing 3D Scalar Fields", Computer Graphics, 22(4): 51-58 may be more efficient, but the results after DL and/or ADL are applied render the substantially identical results.

Referring to Figure 29, the Ray-casting paradigm is based on understanding of the geometric relationships between a model, the screen plane, the user's eye point, and a view direction (expressed in the form of a view vector V), usually perpendicular to the screen plane. For every screen pixel in the projection of the model, we consider the voxels along the "ray" (line) parallel to V which passes through the pixel. In the preferred embodiment of the present invention, in which the model takes the form of a convex

polyhedron, there is associated with every such pixel at screen coordinates (x',y') an "entry point" at coordinates (x',y',z') where the ray enters the model, and an exit point (x',y',z'') where the ray exits the model.

5 Ray casting methods are usually defined as sequential algorithms based on processing the voxels along each ray in sequence from front (i.e., the entry point) to back (i.e., the exit point). The precise means used to determine which voxels are involved are straightforward and known to those of skill in the art of three-dimensional image processing.

10

The previously described texture-mapping technique can be described within the ray casting paradigm as a method where for each ray, the voxel value at the entry point is projected onto the screen pixel, and all other voxels along the ray are ignored.

15

In Maximum Intensity Projection (MIP), the maximum voxel value along the ray is projected onto the screen pixel. An example of MIP is illustrated in Figure 16 for the case of a three-dimensional Doppler ultrasonographic image of blood vessels. MIP is a simple technique which gives some appreciation of
20 three-dimensional structure, but has the disadvantage that it is typically very

difficult to determine occlusion (to tell which of two overlapping structures is in front and which is behind).

A useful variant of MIP is Conditional MIP (CMIP), in which the maximum value of those voxels along the ray which meet some predefined condition are projected onto the screen pixel. In the presently preferred embodiment of the present invention, Doppler ultrasonographic images featuring a mix of gray scale and color information are encoded such that color voxels (representing flowing blood) fall within prespecified value limits. A representative illustration is shown in Figure 17 in which a three-dimensional image and model is produced by texture mapping. This image is from a volumetric image array (color Doppler ultrasound of the carotid artery of a healthy volunteer) containing a mix of colored voxels representing flowing blood and gray scale voxels representing stationary tissue. However, when a CMIP is performed on the image of Figure 17 with the condition that only voxel values in the color range are projected, improved visualization of the three-dimensional structure of the blood vessels is achieved as illustrated in Figure 18. As shown in the present figure, the three-dimensional image was produced using the volumetric image array of Figure 17, but using Conditional MIP to project only the color information. As shown in the Figure, the front face of the model has been translated outward to reveal the entire blood vessel structure.

In Raysum, all voxel values along the ray are summed and the resulting value, which is normally scaled into an acceptable pixel value range by a preset normalizing constant, is projected onto the screen pixel. An example of the Raysum technique is shown in Figure 19 in which the three-dimensional image is generated using the model and volumetric image array of Figure 16. Typically, Raysum yields a more natural-looking result than MIP, however occlusion is still not shown.

In the presently preferred embodiment of the present invention wherein the DL imaging processing means is applied to MIP, CMIP and/or the Raysum techniques, a fixed distance called the "ray depth" is chosen, and the selected process (MIP, CMIP, Raysum or other techniques) ignores voxels beyond the ray depth. The ray depth thereby defines a "thick plane" as a region of interest. In accordance with the Ray-casting paradigm, every screen pixel in the projection of the model, has voxels processed along the "ray" line parallel to V, but only as deep into the model as determined by the ray-depth. This provides a middle ground between methods such as MIP and Raysum, which attempt to display the entire model interior, and texture mapping, which displays only surface detail. Accordingly, the DL technique is particularly useful when examining vascular structures.

In operation, the user selects the ray-depth and process (MIP, CMIP, or Raysum) via a graphical user interface and input device such as a single button mouse. The ray depth can be selected from more than one face when a perspective view is desired. Figures 23 and 24 represent illustrative examples of DL-MIP and DL-Raysum image processing from two faces. Typically, the user selects a face on the model from which an enhanced view is desired. This face is considered as the entry point or starting point from which the ray depth is measured. The user then provides or accepts the default ray-depth which will typically be measured along V perpendicular to the selected model face. The depth determines the exit point within the volumetric array beyond which deeper voxels are excluded from processing. The limiting of voxels from processing in this manner significantly speeds processing time. Furthermore, the processed voxels within the "thick plane" region of interest yields enhanced imaging results which show occlusion.

15

Figures 20a through 20c show the results of applying texture mapping to the same Doppler ultrasonographic data illustrated in Figures 16 and 19 however, the difference between the three figures is that the frontmost model plane is in different positions. As is apparent from the Figures, the complex three-dimensional structure of the blood vessels is not apparent when only planar cross-sections are seen.

20

Figures 21a, 21b, and 21c show the results of applying the DL image processing means to the MIP image of Figure 16 (DL-MIP) to the same data and model as Figure 20a, but varying the ray depth at two millimetres, four millimetres and six millimetres respectively. Figures 21c, 22a and 22b
5 illustrate the effect when the ray depth is maintained at six millimetres and the frontmost model plane is translated inward through the three-dimensional volumetric array to reveal blood vessel structures at different depths.

Figure 23 shows the application of DL-MIP with two model faces placed
10 close within the ray depth of bright structures within the vascular volumetric image data. Figure 24 shows the result when the DL image processing technique is applied with the Raysum technique (DL-Raysum) with the identical volumetric image data and model illustrated in Figure 23.

15 It is important to note that in the preferred embodiment of the present invention, the DL-MIP and DL-Raysum processes can be performed at near real-time speeds, allowing updating of the image while the operator moves model planes on the screen 36a. Movement of the model planes will be described in greater detail below. This interactivity contributes to an enhanced
20 appreciation of the three-dimensional structure by introducing an operator-controlled relationship between depth in the image and time.

Another previously mentioned novel imaging processing means is the Adaptive Depth-Limited technique which may be applied to both MIP and Raysum processed images (ADL-MIP and ADL-Raysum respectively). ADL applies a user selected condition on voxel values, which discriminates between structures of interest and "background" structures in the volumetric image array. The ADL process proceeds in two phases. In the first phase, starting at the entry point, voxels along the ray which do not meet the condition are ignored. In the second phase, Voxels are examined along the ray depth starting at the first voxel which meets the condition, and proceeds in a manner substantially identical to the DL- MIP or DL-Raysum process. If no voxel along the ray meets the condition, the screen pixel is set to a predetermined color, usually black.

The adaptive processes are preferred to the conventional MIP and Raysum techniques in that they can be performed faster as it is not necessary to include every voxel in the computation. Furthermore the ADL image processing means provide a clear indication of occlusion. As an example when comparing Figures 25 and 26, which illustrate ADL-MIP and ADL-Raysum respectively with figures 16 and 19, which illustrate respectively, conventional MIP and Raysum on the same vascular volumetric image data, the ADL processed images clearly indicate occlusion.

It is important to note that in the preferred embodiment of the present invention, the ADL-MIP and ADL-Raysum processes can be performed at interactive speeds. This provides an enhanced appreciation of three-dimensional structure, because the user can freely rotate the image to any desired viewing
5 angle on the screen, with a speed and smoothness which supports the appreciation of parallax depth cues.

The Depth-Limited MIP technique shares some features of the so-called Sliding Thick-Slab MIP or STS-MIP technique described in STS-MIP: A new
10 Reconstruction Technique for CT of the Chest, Sandy Napel, Geoffrey D. Rubin, and R. Brooke Jeffrey, Jr., *Journal of Computer Assisted Tomography* 17(5): 832-838. However, the following key differences should be noted. The STS-MIP technique performs MIP only in a single direction while DL-MIP works in any viewing direction selected interactively by the user. The STS-MIP method
15 is not interactive, and is normally presented in the form of a retrospective cine sequence of two-dimensional images, while in the preferred embodiment of the present invention, DL-MIP is fully interactive and presents image processed three-dimensional images at or near real-time speeds. Finally, STS-MIP applies primarily to X-ray CT angiography of the chest, while in the present invention,
20 DL-MIP is applied to three-dimensional ultrasonographic imaging and other forms of three-dimensional volumetric image data.

The adaptive variants of DL-MIP and DL-Raysum identify for each pixel, a depth value which is the depth of the first voxel which meets the chosen adaptive condition. These depth values can be used to enhance the resulting images by deepening the contrast according to any of several shading methods known to those of skill in the art. Two useful shading methods presently preferred are Depth Shading and Surface Shading.

When Depth Shading is applied, pixel intensities are reduced by a multiplicative factor inversely proportional to depth. In Surface Shading, the collection of depth values for all pixels in the projection of the model are treated as a surface in three dimensions, and the effects of light from one or more light sources reflecting from that surface toward the viewer's eye position are modelled according to methods (e.g. Lambertian reflection modelling) known to those skilled in the art of three-dimensional computer graphics.

15

Accordingly, Figure 28 shows the result of applying surface shading to the result of an ADL-MIP projection which may be compared to Figure 25 which shows an ADL-MIP without shading. Similarly, Figure 27 shows the result of applying surface shading to the result of ADL-Raysum image processing which may be compared with Figure 26 which shows an ADL-Raysum without shading.

Returning now to Figure 7, after the three-dimensional image and the model are displayed on the screen 36a of the monitor 36, the initial view is saved in the memory 82 with the image data array and is indexed as the "Reset" view. Therefore, after initialization, four saved views of the image data array exist and can be retrieved from memory and displayed when an associated icon is selected as will be described. It should be apparent that the number of saved views is arbitrary and that fewer or more saved views may be created and saved. When the initiation is complete, the display module 92 begins to monitor continuously the graphical input device 38 to detect input commands representing desired manipulations to be performed on the displayed image (see Figures 7a to 7d). When input commands are detected by the display module 92, the display module manipulates the displayed image in accordance with the received input commands.

Although the display module 92 has been described as retrieving the two-dimensional ultrasound image data from the physical memory 82 and the other related image data from the memory 88 after the entire target volume has been scanned, it should be appreciated by those of skill in the art that the calibration file may be retrieved by the display module 92 prior to acquisition of the two-dimensional ultrasound image data. In this case, the other related image data in the calibration file does not describe acquired two-dimensional ultrasound image

data per se but rather describes the memory location and characteristics of the structure in physical memory 82 into which the two-dimensional ultrasound images will be captured. When the system is operated in this manner, the display module 92 is able to present a three-dimensional image of the target volume
5 virtually simultaneously as the target volume is being scanned.

Although the system 20 has been described as moving the ultrasound probe 24 along a linear scanning path Z, it should be realized by those of skill in the art that other ultrasound probe scanning geometries may be used. In these
10 instances, the other related image data in the calibration file must be altered to take the scanning geometry into account so that a relatively simple transformation may be used to convert the two-dimensional image data into a volumetric image array.

15 Three-Dimensional Image Manipulation

All manipulations of the displayed image can be executed via three actions using the graphical input device 38. These actions are termed "point", "click" and "drag". To "point" is to move the graphical input device 38 so that the
20 cursor is positioned on the screen 36a at a desired region without depressing its button. To "click" is to press down on the button of the graphical input device

while to "drag" is to move the graphical input device while the button is depressed. The term "point-click-drag sequence" is used to denote the complete action of moving the cursor into a desired region on the screen via the graphical input device, depressing the button on the graphical input device, moving the graphical input device to another desired region with the button depressed and then releasing the button. Specific manipulations of the image are carried out according to the region in which the click aspect of the point-click-drag sequences occur.

10 Once the initialization routine has been completed and the model and three-dimensional image are displayed on the screen, the graphical input device is monitored to determine whether the user wishes to manipulate the displayed image. The manipulations supported by the display module 92 are rotation of the entire model and three-dimensional image about an arbitrary axis, translation of
15 a selected plane of the model and rotation of a selected plane of the model about an arbitrary axis. The manner in which the display module 92 interprets movement of the graphical input device 38 and manipulates the displayed model and image in response to the graphical input device 38 will now be described.

20 The display module 92 continuously monitors the graphical input device 38 to determine the position of the cursor on the screen and to determine if a

click has occurred. Specifically, the display module 92 determines if the cursor is located within the exterior area of a visible model face (block 324). If the cursor is positioned within such an exterior area, the display module 92 colors the perimeter lines of that model face blue (block 326). Otherwise, the display
5 module 92 changes the color of any perimeter lines which may have been blue to their previous color (block 328). The display module 92 also determines if a click has occurred (block 330). If a click has not occurred, the movement of the graphical input device 38 is tracked. The display module 92 moves the cursor on the screen 36a to follow the movement of the graphical input device (blocks
10 332 and 334). As the cursor is moved, the display module 92 reverts back to block 324 so that the perimeter lines of the visible faces change colors depending on the location of the cursor as just described.

When a mouse click is detected, the position of the cursor on the screen
15 36a is examined to determine if the cursor is located within the main display window (block 336). If the cursor is outside of the main display window, the display module 92 determines if the cursor is positioned over an option icon (block 338). If the cursor is not over an option icon, the click is ignored and the graphical input device 38 is monitored until the button is released (block 339).
20 At that time, the display module 92 reverts back to block 324. If however, the

cursor is positioned over an option icon, then the display module 92 executes a routine associated with the selected option icon as will be described.

When a graphical input device click occurs and the cursor is positioned
5 within the main display window as detected at block 336, the display module 92 determines whether the cursor is located within the interior area or exterior area of a displayed model face or in the background (block 340). If the cursor is located in the background, the display module 92 determines that the user wishes to rotate the entire model and three-dimensional image. In this case, after the
10 button has been depressed, the drag direction and drag distance of the graphical input device 38 is monitored (block 342). As the graphical input device 38 is being dragged, the drag direction and drag distance are repeatedly conveyed to the display module 92 (block 344). The drag distance and drag direction values are used by the display module 92 to rotate the three-dimensional model and
15 image and to update the display to show the rotation on the screen (block 345). Once the button on the graphical input device 38 has been released, the manipulation is considered complete and the display module 92 reverts to block 324 (block 346).

20 A variety of techniques for converting mouse drag distance and direction to three-dimensional rotation axis and angle parameters are known to those

skilled in the art. The preferred embodiment of the present invention uses Shoemake's "Arcball" technique described in the Proceedings of Graphics Interface 92 published by the Association For Computing Machinery (ACM). The fixed point of rotation for the three-dimensional image is constrained to be
5 the geometric center of the initial model. Thus, in this manipulation, movement of the graphical input device vertically on the screen 36a through the center of the displayed image causes the displayed image to rotate about a horizontal axis, while movement of the graphical input device horizontally through the center 34 of the image causes the displayed image to rotate about a vertical axis.

10

Figures 8a to 8c show the model and three-dimensional image within the main window display undergoing a rotation about a vertical axis as the graphical input device 38 is moved to drag the cursor across the main display window from mid-right to mid-left. Figures 9a to 9c show the model and three-
15 dimensional image undergoing a rotation about an axis, angled at about 30° to the horizontal and sloping up and to the right, as the graphical input device 38 is moved to drag the cursor across the main display window from top-left to bottom-right. As should be apparent, this operation gives the user the sense of taking hold of the displayed image and pulling it around. Further details of this
20 image manipulation can be found in the above-mentioned publication.

To facilitate understanding of manipulations of the model other than simple rotation of the entire model, it is necessary to describe the polyhedral model in greater detail. Mathematically, a convex polyhedron can be characterized as the intersection of a set of half-spaces defined by at least four
5 planes, herein called bounding planes. Each face of the polyhedron is a convex polygon embedded in the corresponding bounding plane. By changing the parameters of the bounding planes (i.e. the coefficients the display module 92 either to retrieve a copy the sheared volumetric image array from the memory
88 A,B,C,D in the plane equation $Ax + By + Cz = D$), the shape of the model
10 polyhedron can be modified. The number of bounding planes may also be changed. Specifically, new bounding planes may be added and existing planes removed from the mathematical specification of the model. The result is that the model polyhedron gains or loses faces.

15 The display module 92 supports two primary manipulations of bounding plane coefficients, namely translation (change of coefficient D, which essentially specifies the perpendicular distance from the plane to the coordinate origin) and rotation (change of coefficients A,B,C, which collectively specify the orientation of the plane relative to the coordinate axes). As will be described below, the
20 choice of which bounding plane (and hence which corresponding model face) is to be affected, and whether to perform translation or rotation, is determined by

contextual interpretation of point-click-drag sequences relative to the displayed model. The display module 92 also provides means to add and delete bounding planes from the model specification, as will also be described below.

5 A distinction is made between original bounding planes, which are aspects of the model assigned to the volumetric image array $V(x,y,z)$ when it is first loaded into memory (blocks 310 and 312 in Figure 6) and planes added in response to user input. Model faces corresponding to original bounding planes have their perimeter lines displayed as white lines, while faces corresponding to
10 added planes are indicated using another color, typically yellow or green. Only added planes may be translated, rotated or deleted. The original planes represent the boundaries of the volumetric image and, provide the means to support the addition of new planes.

15 If at block 340, the click is detected and the cursor is determined to be within the perimeter of a displayed model face, the face of the model in which the cursor is positioned is determined and the position of the cursor within the face is examined by the display module 92 (block 350 in Figure 7c). If the
20 cursor is positioned within the interior area of the face, the display module 92 determines that it is desired to translate the corresponding bounding plane. Thereafter, the display module 92 examines the corresponding plane to

determine whether the plane is an original plane i.e. one denoted by white lines (block 354). If the plane is an original plane, a new plane is created and added to the model (block 356). Initially, the added plane is congruent to the original bounding plane. Once the added plane has been created or if at block 354, the
5 plane is not an original plane, the perimeter lines denoting the plane (i.e. the perimeter of the corresponding displayed model face) are colored yellow and all perimeter lines of other added planes are colored green (block 357).

The drag distance and direction of the graphical input device are then
10 monitored (block 358). The display module 92 next determines the direction of translation of the added plane. To do this, the display module 92 calculates the dot product of the drag vector and the projection onto the screen 36a of a normal vector to the plane is computed. If the dot product is positive, the plane is translated in the direction of the normal vector; if negative, it is translated in the
15 opposite direction. In the preferred embodiment, the model specifications are such that all boundary plane normal vectors point away from the interior of the model polyhedron. Hence, movement of the graphical input device 38 which results in a positive dot product pulls the plane outward from the center of the model, while movement of the graphical input device 38 which results in a
20 negative dot product pushes it in (block 360).

If the translation direction of the plane is determined to be in, the display module 92 checks to ensure the translation does not result in the collapse of the model (block 362). If it does, the display model 92 reverts back to block 358 without updating the displayed model and three-dimensional image. Otherwise, translation of the added plane occurs, and the translation of the added plane occurs, and the image is updated to reflect the changed polyhedral model, i.e. with the texture mapping technique. The points in the volumetric image array $V(x,y,z)$ which correspond to those on the translated plane are texture-mapped onto the plane (block 364).

10

If a block 360, the translation direction is determined to be out, the display module 92 checks to see if the translation can result in the plane disappearing (block 366). If not, the display module 92 proceeds to block 364 and updates the displayed model and three-dimensional image. Otherwise, the display module 92 notes that the translation may result in the disappearance of the plane before proceeding to block 364 (block 368). As should be apparent, as a plane is translated, various cross-sections of the image can be viewed as the translated plane slices through the volumetric image array $V(x,y,z)$. Once the plane has been translated, the display module checks to see if the plane was deleted (block 370). If not, the perimeter lines of the translated plane are colored yellow and all other perimeter lines are colored green (block 372).

20

Otherwise, the perimeter lines of the next most recently modified plane are colored yellow and all other perimeter lines are colored green (block 374). When the button is released, the display module 92 considers the translation manipulation complete and reverts back to block 324 (block 376).

5

When the added plane which is to be translated is nearly parallel to the plane of the screen 36a, the dot product described above is essentially zero. In such cases, the display module 92 only considers the vertical component of the graphical input device movement to determine the translation direction and distance. In this instance, upward movement of the graphical input device 38 causes the display module 92 to push the added plane into the model 36a while downward movement of the graphical input device causes the display module to pull the plane out of the model. Figures 10a to 10c show the model and image within the main display window, wherein a plane of the model is translated towards the geometric center of the model.

10
15

When a click is detected and the cursor is positioned within the exterior area of a model face as determined at block 350, the display module 92 determines that the corresponding bounding plane is to be rotated. The display module then examines the plane to determine whether the plane is an original plane (block 380). If the plane is an original plane, a new plane congruent to the

20

original plane is created and added to the model (block 382). The perimeter lines of the added plane are colored yellow. Once this has been done or if at block 380, the plane is not an original plane, the perimeter lines of all other added planes are colored green (block 384).

5

After this, the drag distance and direction of the graphical input device 38 are monitored and the rotation axis and direction of the graphical input device are computed by the display module 92 using Shoemake's technique referred to previously (block 386). After this, the display module 92 determines whether
10 the rotation would cause the plane to disappear (block 388). If so, the display module 92 reverts back to block 386 without updating the displayed model and three-dimensional image. Otherwise, the display module 92 rotates the added plane about the geometric center of the initial model by the calculated amount. As the rotation is occurring, the display module 92 updates the image on the
15 screen (block 390). This allows different cross-sections of the image to be viewed as the rotated plane slices through the volumetric image array $V(x,y,z)$. Figures 11a to 11c show the model and three-dimensional image within the main display window, wherein a plane of the model is rotated about an axis, angled at about 30° to the horizontal and sloping up and to the right. It is through this
20 manipulation that new oblique planes may be added to the model. Once the button on the graphical input device 38 has been released signifying that the

plane rotation has been completed, the display module 92 reverts to block 324 (block 392).

As should be apparent from the above description, whenever original
5 planes of the model are selected by the user to be translated or rotated, added planes are created and it is the added planes which are moved. The original planes, although not shown on the display screen, remain stored in memory 82 and may be recalled at any time to reset the display to its original state. It should also be apparent that white lines denote an original plane, green lines
10 denote an added plane with the exception of the last moved plane which is denoted by yellow lines and blue lines denote that the plane is going to rotate if a mouse-drag sequence occurs.

When an added plane is rotated, the added plane may become oblique but
15 it is not permitted to extend beyond the boundary defined by the original planes (see Figures 11a to 11c). Any added plane can be translated sufficiently far out of the model that its corresponding model face disappears. If the user releases the graphical input device button after the face has disappeared, the plane is removed from the model. This allows a user to delete unwanted added planes.
20 If the graphical input device button is not released, the user can push the added plane back into the model so that the plane becomes visible again and no deletion

takes place. Figures 12a to 12d show the model and three-dimensional image within the main window display wherein an oblique plane of the model is translated away from the geometric center of the model until it disappears. Although it would seem that if an added plane is translated sufficiently into the
5 model, the entire model can collapse, the display module 92 does not permit an added plane to be translated so far as to collapse the model (see block 362).

While the graphical input device 38 is being moved to effect changes in the displayed view and the display is updated showing intermediate positions and
10 orientations of the affected plane or planes, the display module 92 must re-sample the volumetric image array $V(x,y,z)$ and complete the rendering process, a process of discrete approximation. In this embodiment, a number of re-sampling methods are available, each of which offers a different compromise between computational speed and image quality.

15

To achieve smooth operation with limited computer power, the display may be computed (rendered) at less than the full resolution of the monitor screen and/or a simpler interpolation technique may be employed in the re-sampling process. In this embodiment, the display is computed in as many as three
20 rendering passes, the first and third of which may be disabled if the user so wishes, by selecting the appropriate option icon via the graphical input device

38. The enabled/disabled status of each pass is what is actually set to a default state at block 324 during initialization. The first enabled pass in the sequence is uninterruptible, i.e. while the graphical input device is being moved, the first enabled pass is performed in its entirety, yielding a succession of complete views
5 on the screen. Subsequent enabled passes are automatically interrupted by graphical input device movement, the visible result being that the displayed view is only replaced by a higher quality view (computed by an interruptible rendering pass) when there is a sufficient pause in graphical input device movement. The three rendering passes supported by the present embodiment
10 are:

1. reduced image resolution, nearest-neighbour re-sampling
2. full image resolution, nearest-neighbour re-sampling
15
3. full image resolution, tri-linear interpolation re-sampling

As mentioned previously, at block 338, if a click is detected and the cursor is not positioned in the main display window, the display module 92
20 determines whether an option icon has been selected. The available option icons allow the user to select parameters different from the default values, to enhance

image display and to execute special feature routines. These option icons include "Reset", "Views A to C", "Remember", "Snapshot", "Animation", "Indicator", "Orientation", "Fast", "Smooth", "Win", "UV", "Magnify" and "Measure". Figure 13 illustrates most of these option icons in a control display window positioned beside the main display window. In this example, the preferred Views A to C have been labelled by the user as "Sagittal", "Coronal" and "Axial". The available options which can be selected via an option icon will now be described.

10 If the Reset icon is selected, the original view of the image and model stored with the volumetric image array $V(x,y,z)$ is recomputed and displayed on the screen. Likewise, if one of View A to C icons is selected, the corresponding preferred view is recomputed and displayed. If the user wishes to change one or more of the Views A to C for a current session, the user can substitute the displayed view for the stored view. The present embodiment permits the user to activate a distinct window in which the View icon labels (eg. Sagittal, Coronal, Axial etc. in Figure 13) are displayed and to edit the labels as desired. Changes made to the labels persist only for the current session, unless the user elects to save the changes in memory 88 using an option icon provided for that purpose, in which case any preferred Views associated with the data file are overwritten.

If the Remember icon is selected, the current view on the screen is stored in memory 82 overwriting the "Reset" view for the current session only. The "Reset" view associated with the current data file in memory 88 is not changed, only the copy in memory 82. This view may be recalled to the screen at any
5 time by selecting the Reset icon, unless and until it is overwritten by a subsequent use of the Remember icon.

It should be realized that a similar two-icon technique can be used for Preferred Views A to C. However, the present embodiment allows the user to
10 overwrite these views in memory 82 by holding down a specified key on the keyboard while selecting the corresponding View icon.

If the Snapshot icon is selected at any time during manipulation of the model and image, the image currently displayed in the main display window is
15 stored as a file in memory 88 in an industry-standard image file format, in order that it be may subsequently be used with other software. The present embodiment uses a tagged image file format ("TIFF"). It should be realized that adding support for other file formats can be achieved in a straightforward manner by following published format specifications.

20

If the Animation icon is selected, animated sequences of displayed views can be created and saved into memory 82 in an industry-standard image format as just described. When the Animation icon is selected, the display module 92 determines whether a view of the image has been saved using the Remember icon and retrieves it. If no view has been saved using the Remember icon, the original view is retrieved. While this is occurring, an animation display window appears on the screen 36a. The display window allows the user to select the number of intermediate views of the displayed image which are to be computed and displayed (see Figure 14). The animation display window also allows the user to adjust the image size, assign an identifier to the animation sequence and preview the animation to ensure that the selected parameters are satisfactory. After this, the display module 92 computes the view orientation, position and orientation of each plane of the intermediate views by simultaneous interpolation between the saved and current views.

15

By using simultaneous interpolation, the user need only enter two views making the use of the feature very simple. Secondly it allows complicated view sequences to be produced which cannot be produced manually. When altering an image manually, a plane may be either rotated or translated but not translated and rotated at the same time. Simultaneous interpolation of plane position and orientation makes it possible to produce an animated view sequence in which a

20

plane is rotated and translated at the same time. As should be realized, this feature as described can only be implemented when the current and saved views have the same number of planes.

5 If the Indicator icon is selected, the model is displayed with axis indicators to indicate standard directions of the image, such as front F, back B, left L etc. These symbols may be changed by the user to suit the application. For example, in ophthalmic imaging the symbols can represent the standard ocular axes, namely superior S, inferior I, nasal N and temporal T. These
10 indicators float as the view of the displayed image changes. To avoid cluttering, it is preferred that axis indicators disappear when the current view orientation would place them behind the displayed model.

In some applications such as in medical imaging, it is desired to show the
15 current spatial relationship of the model to a representation of the structure which has been imaged. This can be achieved when the Orientation icon is selected. When this icon is selected, the representation of the structure being imaged is selected from a list of stored structures. The structure is modelled using conventional three-dimensional computer graphics techniques. The
20 structure is then displayed as a semitransparent solid object intersecting the model polyhedron, either in the main display window or in a second display

window elsewhere on the monitor screen. This requires use of a rendering process which makes it apparent how the structure and model polyhedron intersect. The position, size and spatial orientation of the structure relative to the model polyhedron, which may be expressed as a 4 x 4 transformation matrix, must be determined. When this icon is selected, the structure display window is updated as the user manipulates the displayed view, so that the two displays are always oriented in the same manner. When computer power is limited, it is permissible to update the structure display less frequently than the main display window, e.g. to suppress updating the former until there is a pause in user input. This approach, which works best when the structure is displayed in a window distinct from the main window, is used in the present embodiment. Figures 15a to 15c show the model and three-dimensional image within the main window display as well as the structure of an eye in a second display window beside the main window display. As can be seen, in Figures 15a and 15b, the second display window is small and is positioned below the control display window while in Figure 15c, the second display window has been increased in size.

The Fast and Smooth icons may be individually selected or deselected to enable or disable the first and third rendering passes described previously (the second pass is always enabled). The initial state of these icons is established

during initialization at block 324. It should be realized that this general scheme can be altered slightly, e.g. by addition of a fourth pass with a corresponding option icon to selectively enable or disable it if desired.

5 Each displayed point of the image array $V(x,y,z)$ is converted to a pixel brightness or color by pseudo-color mapping. The domain of the pseudo-color mapping is the range of values in the volumetric image array $V(x,y,z)$. The pseudocolor mapping can be adjusted by a user via the window and level slide controls (labelled "Win" and "Lev" in Figures 13 to 15) to allow the contrast,
10 brightness etc. of the display to be enhanced. The terms "window" and "level" and their interpretation have become standardized in the medical imaging field. The present embodiment is consistent with established medical imaging practice in this regard.

15 When the Magnify icon is selected, a magnifying window appears superimposed upon the main display window and can be moved over the displayed view. Cross-hairs are located at the center of the window and can be positioned over a certain area of the displayed view. When the cross-hairs are at the appropriate location, the user can use the graphical input device to adjust
20 the magnification of the area at which the cross-hairs are located.

When the Measure icon is selected, a measure display window appears on the screen (see Figure 14). The user can use the graphical input device to measure distances and areas of the three-dimensional image within the most recently moved plane (i.e. the one denoted by yellow lines). If the user wishes
5 to measure a distance, the user simply needs to use the graphical input device 38 to indicate the two end points over which the distance is to be measured. If an area is to be measured, the user must identify at least three points. When the cursor is moved over the most recently moved plane, it changes to cross-hairs to facilitate placement of the points on the image. The display module 92 in this
10 mode connects adjacent points by straight line segments and computes both the overall line length and the area bounded by the lines joining the points using an appropriate scale.

If it is desired to change the drag sensitivity, the user can depress a
15 specified key on the keyboard and this will adjust the scaling factor for as long as the key is depressed. This allows for greater or smaller mouse movements for a given change in the displayed image. This of course can also be achieved by increasing or decreasing the size of the main display window.

20 Although the animation function has been described to animate a sequence of translated and/or rotated images, the cine loop technique may also be applied

to visualize four dimensional image arrays. In this case, the four-dimensional image arrays are a series of three-dimensional images of the same target volume, acquired at different times. For example in trans-oesophageal cardiac ultrasound imaging, it is possible to capture and reconstruct a plurality of images of the
5 beating heart, each corresponding to a different point in the cardiac cycle. The three-dimensional images will all have the same spatial interpretation and hence, it is possible to impose a single bounding model upon all of them simultaneously. The intersection of the model polyhedron with each volume image yields a two-dimensional image. The images can be computed and
10 displayed in time order as a cine loop.

Although the system has been described as including a single button mouse to allow a user to input commands, it should be apparent to those of skill in the art that other input/output devices such as a multi-button mouse, a
15 digitizer, a light pen, a trackball, a keyboard or the like or any combination of the above can be used. When other input/output devices are used, different inputs can be chosen to represent different commands or to select the various option icons.

20 When a graphical input device other than the single button mouse is used, manipulation of an image by rotation of a plane of the model can be enhanced.

For example, if a mouse is used in combination with a keyboard and a keyboard input is used to signify that it is desired to rotate a plane, the fixed point of rotation of the plane can be determined by the position of the mouse when it is clicked and the keyboard input signifying rotation of the plane is selected. This
5 allows the plane to be rotated about a point different from the geometric center of the initial model.

If a two button mouse is used as the input/output device, one of the buttons can be assigned to signify translation of a plane and the other button can
10 be used to signify rotation of a plane. This avoids the need to examine the position of the cursor to determine whether it is within interior or exterior areas of the selected model face.

Although the three-dimensional image display technique has been
15 described in an ultrasound three-dimensional imaging system, it should be apparent that the image display technique may be used in different environments where a three-dimensional image needs to be manipulated to allow different views of the image to be visualized. Specifically, the present display technique can be used to display any volumetric image array $V(x,y,z)$ which represents a
20 discretely sampled function defined over a three-dimensional space. This allows

the display technique to be used in other environments such as, for example, magnetic resonance imaging (MRI) and x-ray computed tomography (CT).

Although the system 20 has been described as including a clinical
5 ultrasound machine 28 and a computer 32, it is contemplated that a single
machine be used to perform the functions of both of these components.

Variations and modifications of the present invention should be
apparent to those of skill in the art without departing from the scope of the
10 present invention as defined by the dependent claims.

We claim:

1. A three-dimensional imaging system for acquiring a succession of two-
5 dimensional images of a target volume represented by an array of pixels
 $I(x,y,z)$ into a three-dimensional image represented by a volumetric image
array $V(x,y,z)$ comprising:

scanning means to scan said target volume along a predetermined
geometric scanning path and generate a succession of digitized two-
10 dimensional images thereof representing cross-sections of said target volume
on a plurality of planes spaced along said scanning path;

memory storing said succession of digitized two-dimensional images
together with other related image data defining the location of said two-
dimensional images in said memory and defining interpretation information
15 relating to the relative position of pixels within said two-dimensional images
and to the relative position of pixels in adjacent two-dimensional images
within said target volume;

transformation means receiving said digital two-dimensional images
and said other related image data and transforming said two-dimensional
20 images and said other related image data into a volumetric image array;

display means to generate a three-dimensional image of said target volume from said volumetric image array; and

an image processing means for processing said three-dimensional image of said target volume to render an enhanced three-dimensional image.

5

2. An imaging system as defined in claim 1, wherein said geometric scanning path is linear.

10 3. An imaging system as defined in claim 2, wherein said cross-sections are tilted with respect to an axis normal to the linear scanning path and wherein said linear transformation means transforms said two-dimensional images and said other related image data into said volumetric image array using a shear transformation.

15

4. An imaging system as defined in claim 3, wherein said other related image data includes an address pointer indicating the location of said memory at which the two-dimensional image data begins together with data representing the number of pixels along x and y axis of each two-dimensional
20 image.

6. An imaging system as defined in claim 4, wherein said other related image data further includes physical distance values between the centers of adjacent pixels of the two-dimensional images in the x and y directions and the physical distances between corresponding pixels in adjacent two-
- 5 dimensional images as well as said tilt angle.
7. An imaging system as defined in claim 1 wherein the image processing means employs a predetermined ray depth, thereby defining a "thick plane" region of interest from at least one face of a model within said volumetric
- 10 array and within which voxels are processed.
8. An imaging system according to claim 7 wherein the voxels within the thick plane region of interest are processed by a user selected processing technique.
- 15
9. An imaging system according to claim 8 wherein the user selected processing technique is Maximum Intensity Projection.
10. An imaging system according to claim 8 wherein the user selected
- 20 processing technique is Conditional Maximum Intensity Projection.

11. An imaging system according to claim 8 wherein the user selected processing technique is Ray Summation.

12. An imaging system according to claim 8 wherein the users specifies
5 the ray depth via an input means determined from a model face toward the target volume.





Fig. 1






Fig. 3




Fig. 3a

5/45




Fig. 3b




Fig. 3c

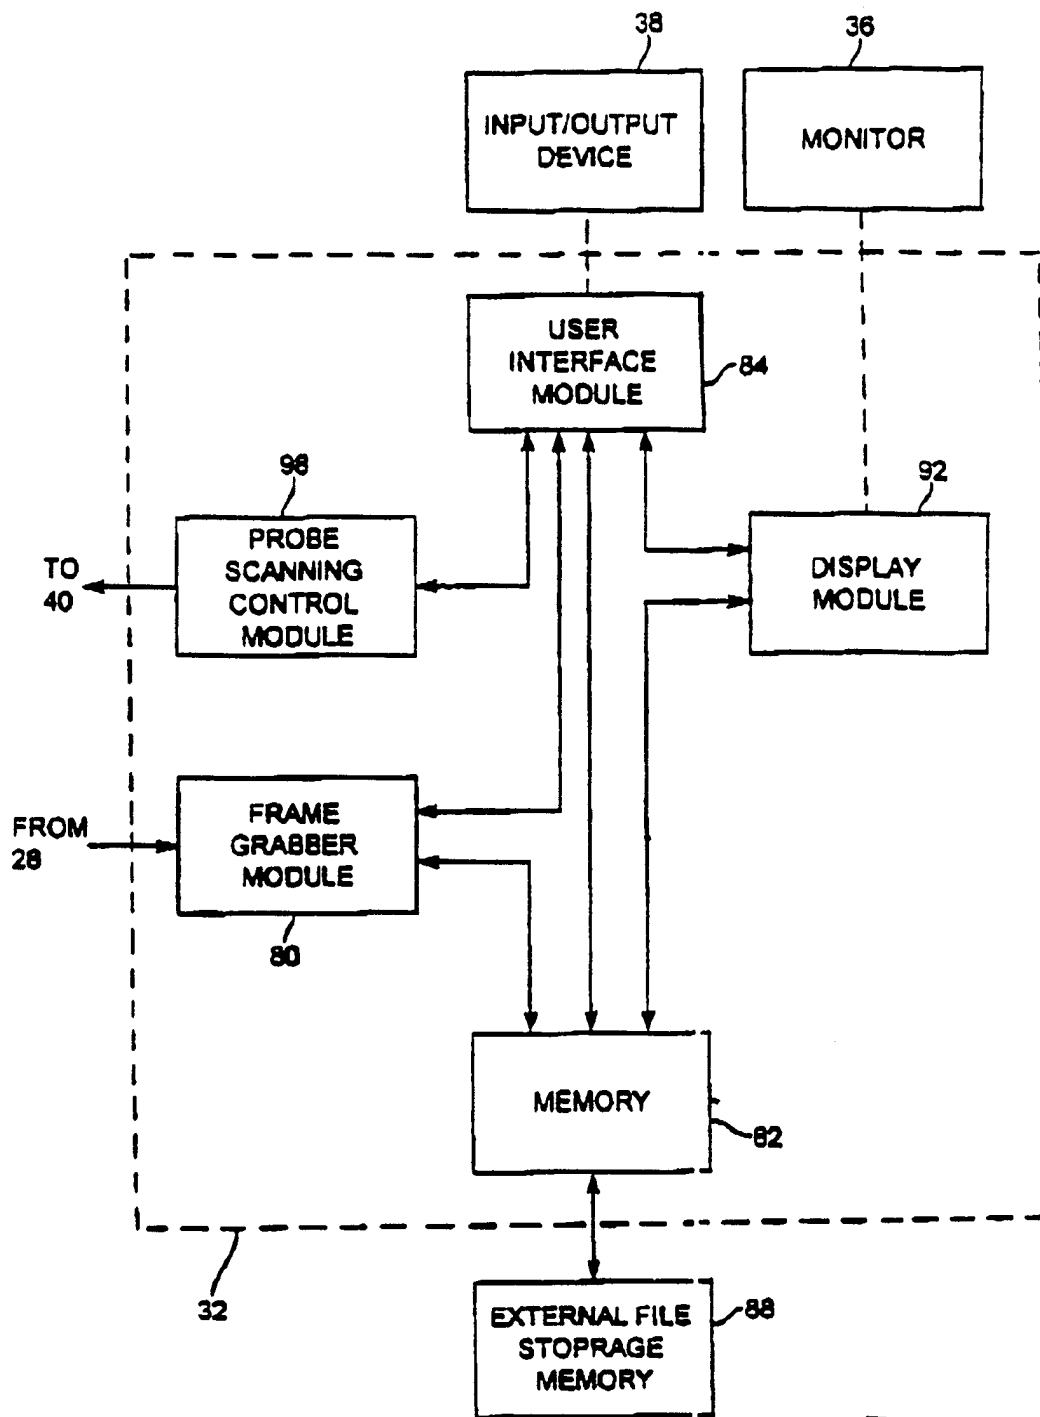


FIG. 4




FIG. 5a

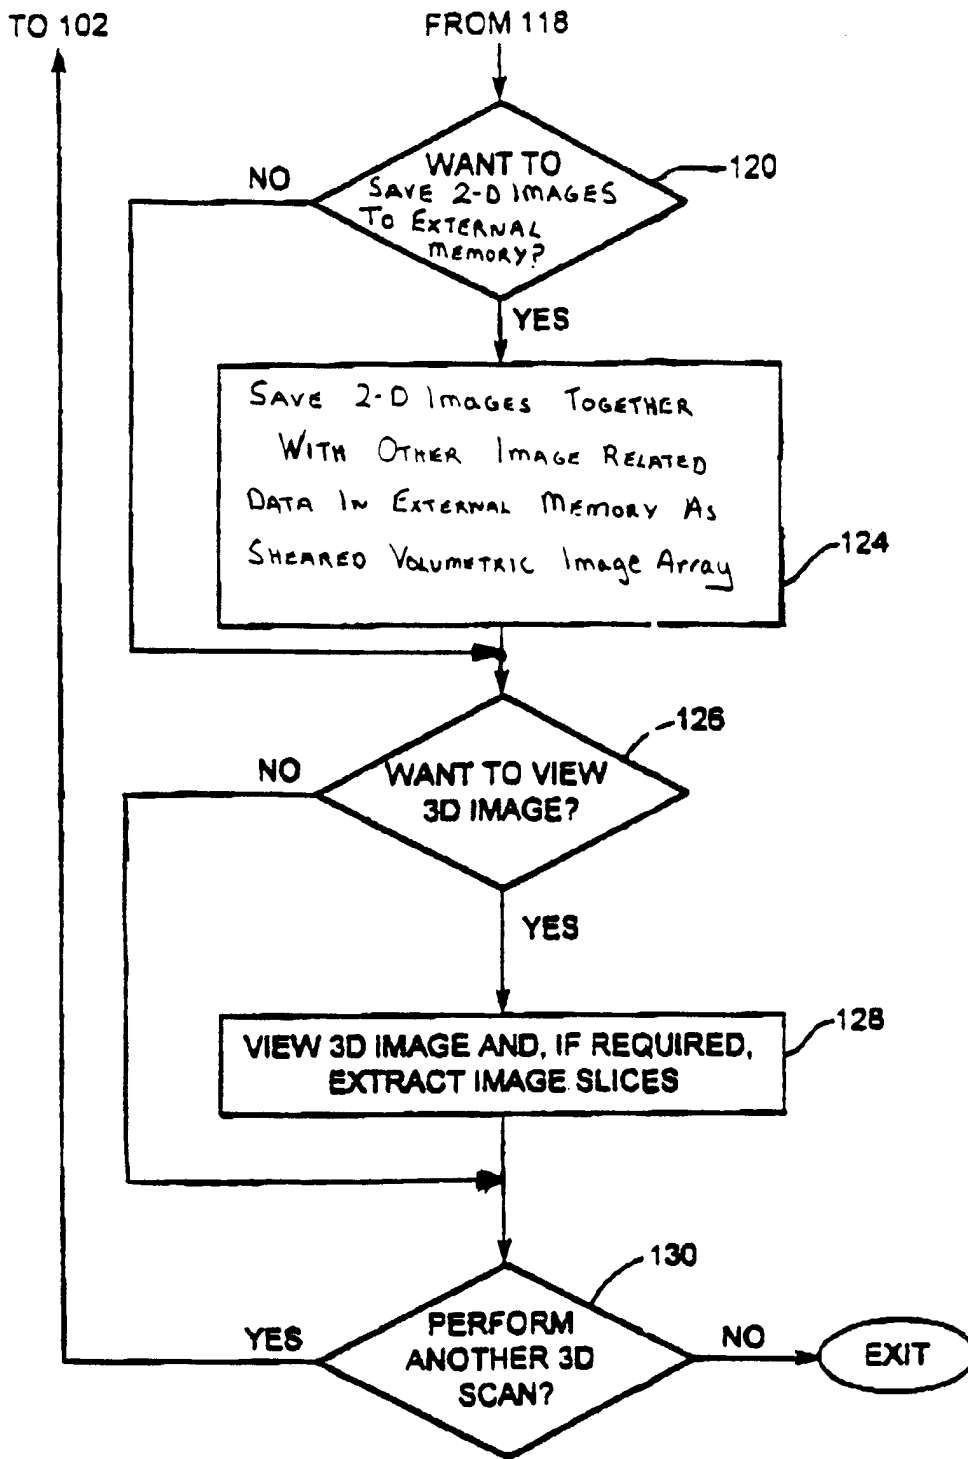


FIG.5b

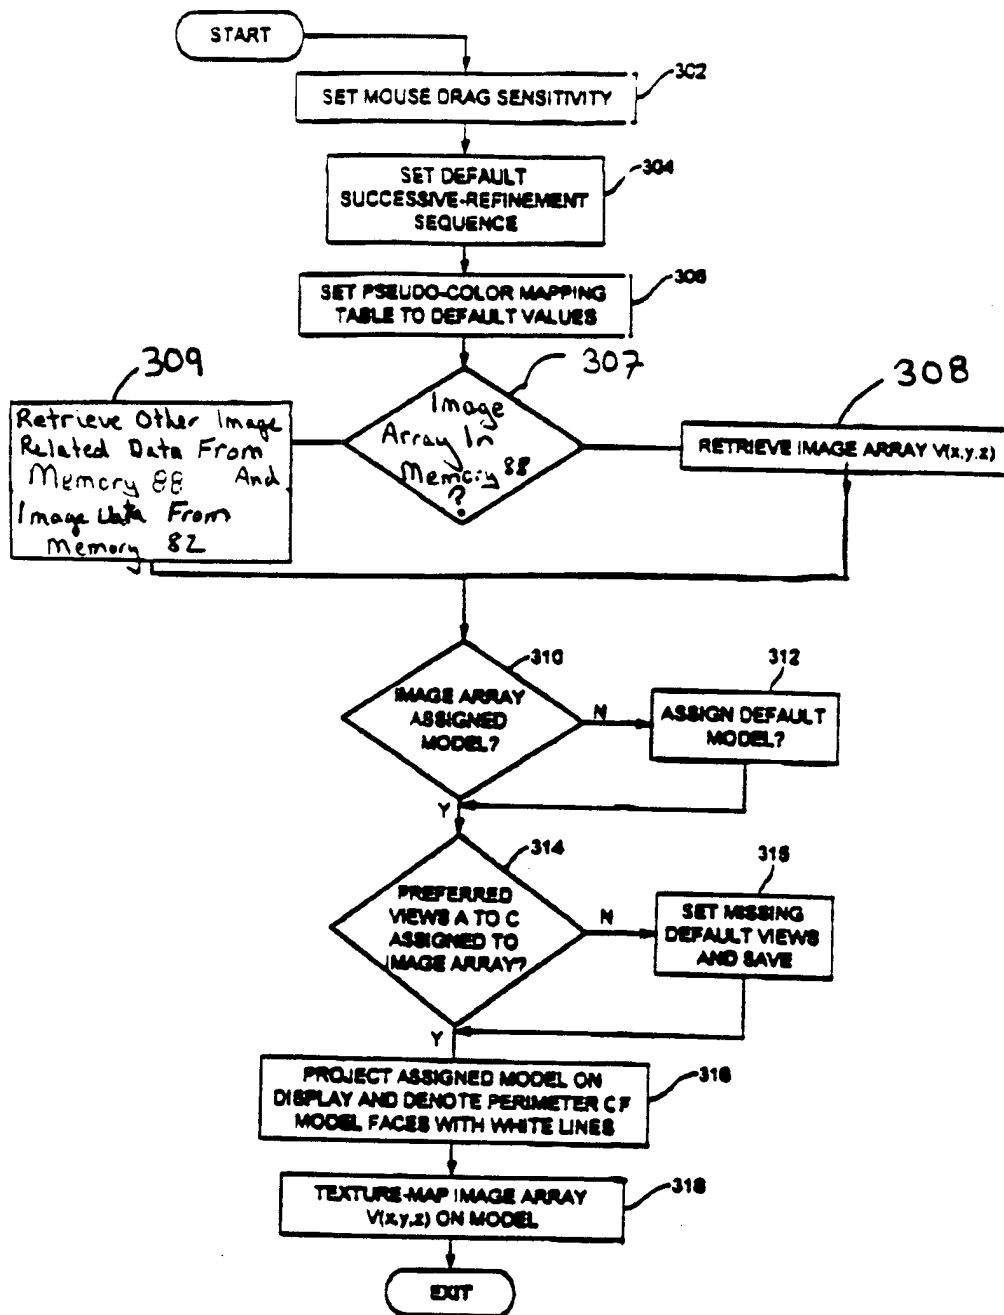


FIG. 6

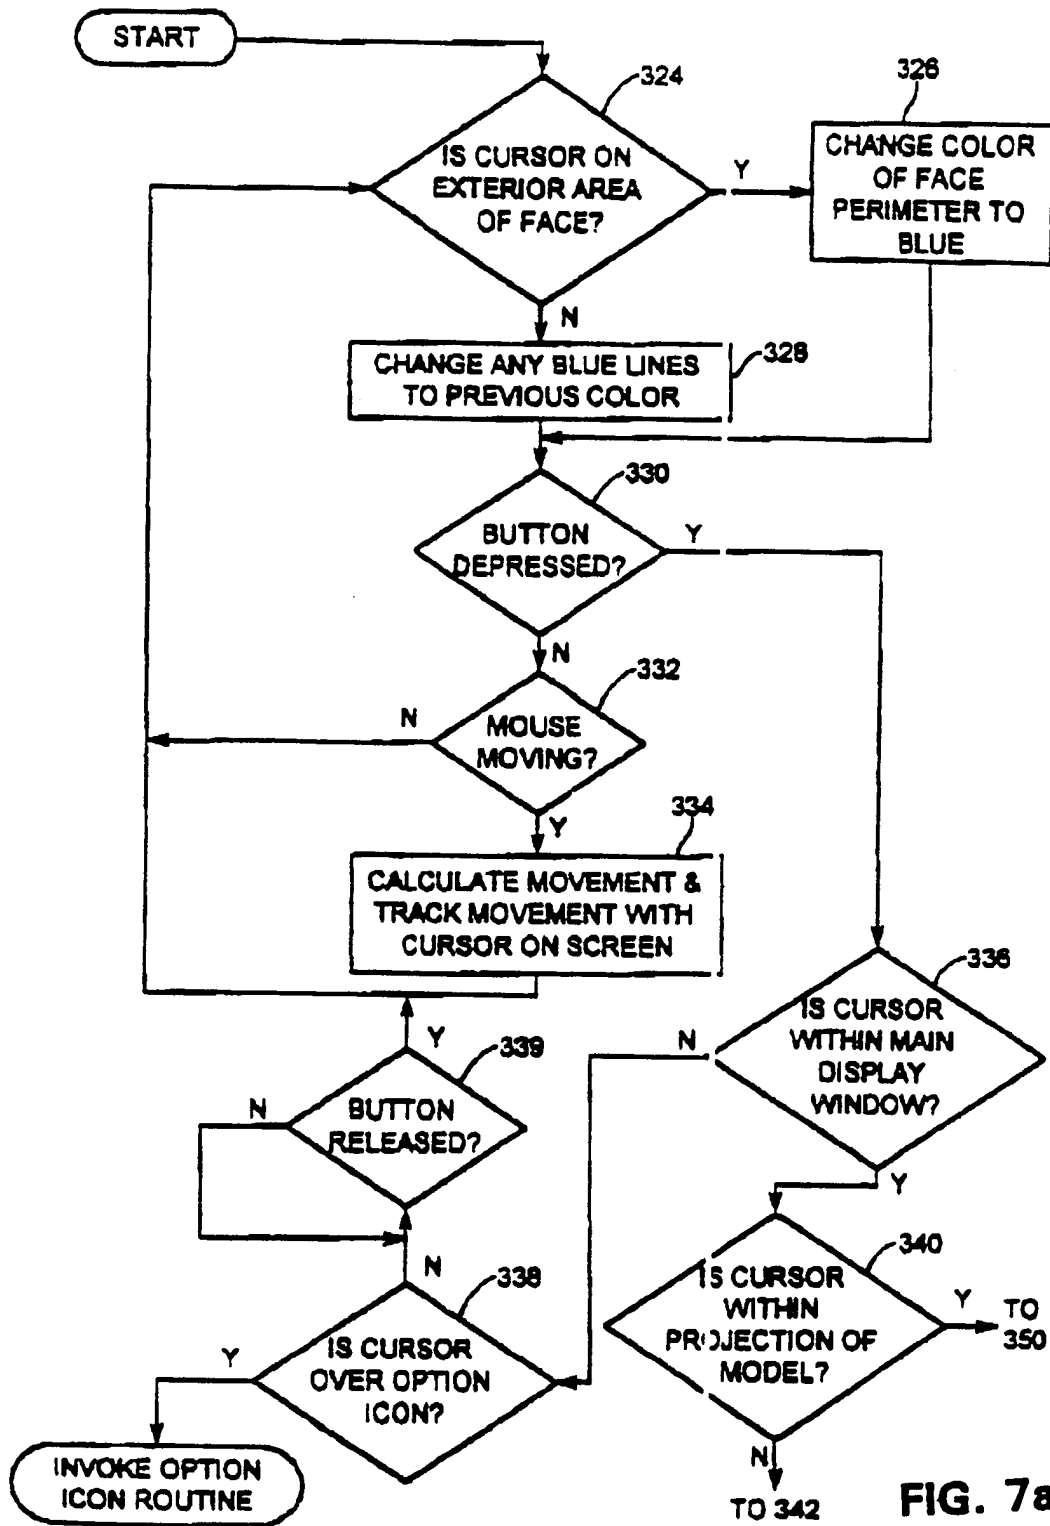


FIG. 7a




FIG. 7b

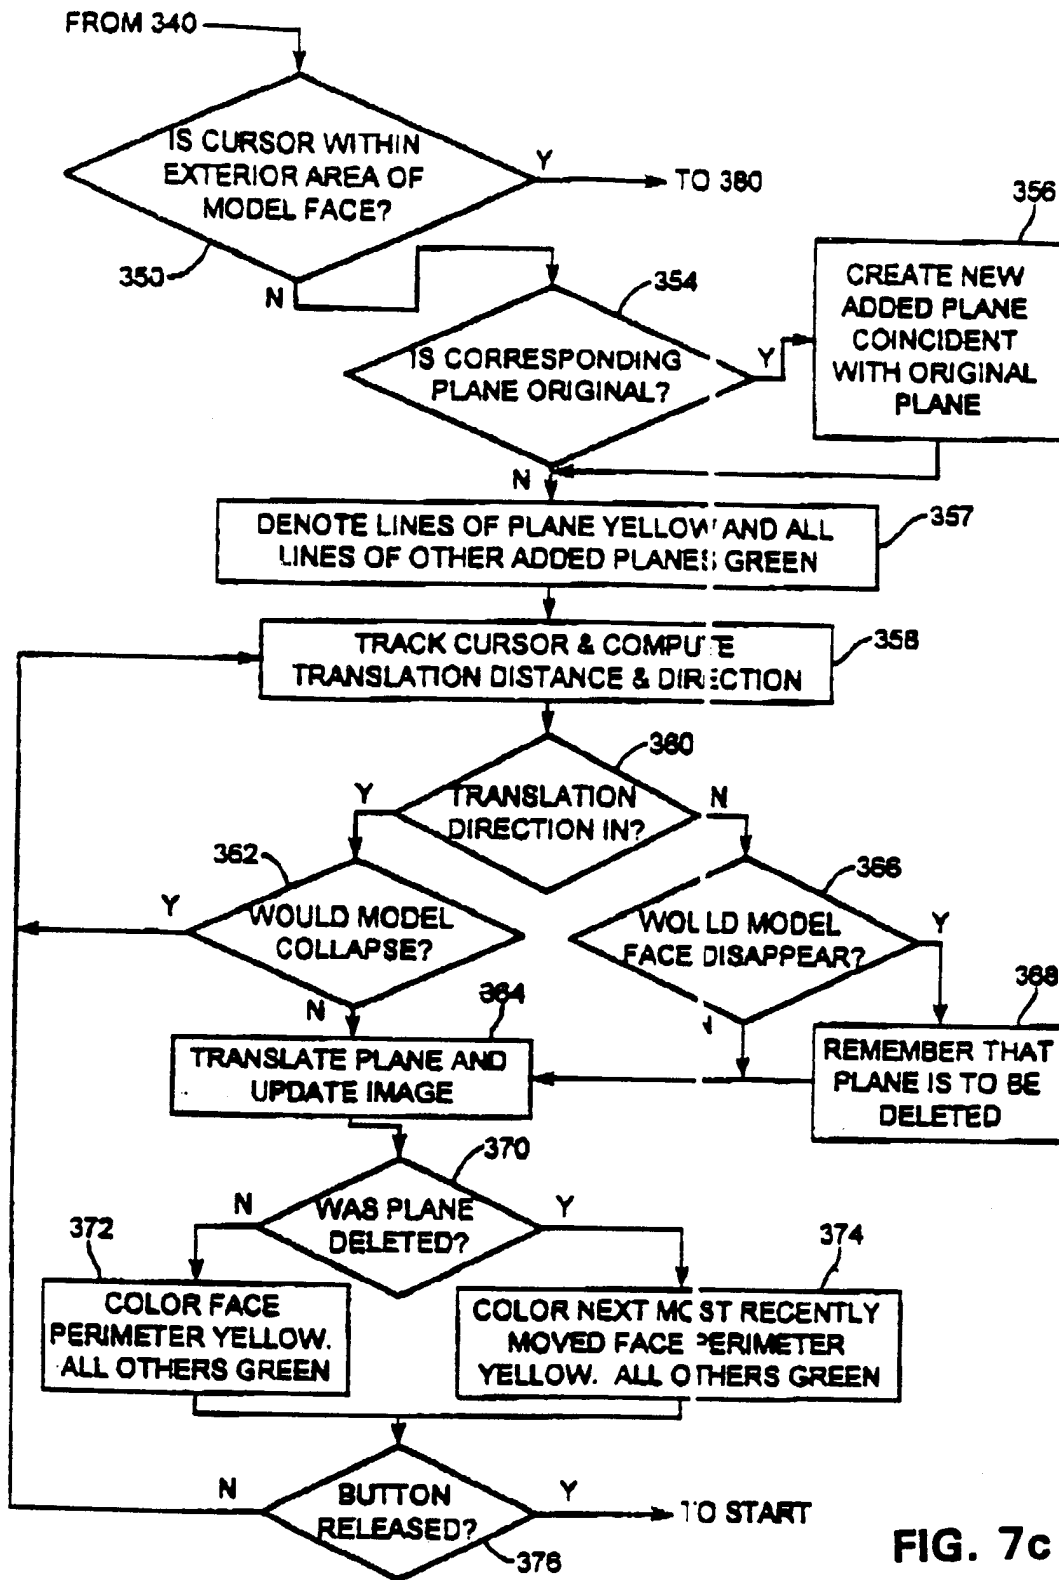


FIG. 7c

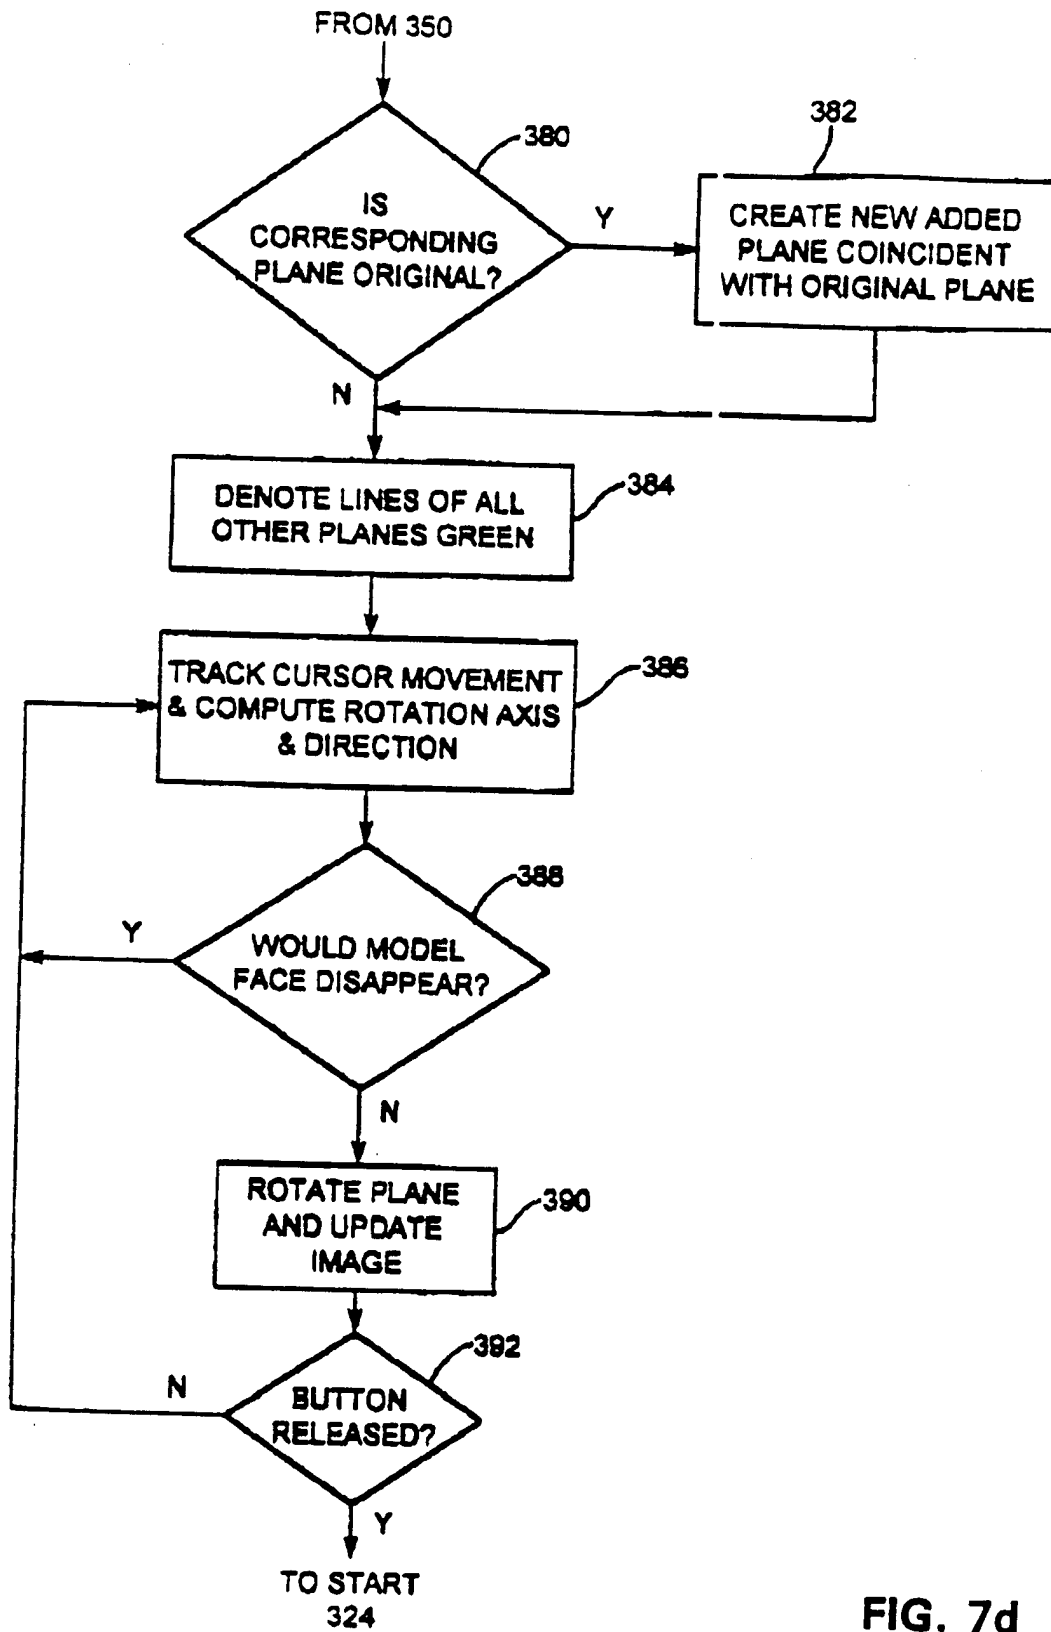


FIG. 7d




FIG. 21a 8a




FIG. 21b 8b

15/45




FIG. 21e 8c




FIG. 22a 9a




FIG. 22b 9b

17/45




FIG. ~~22c~~ 9c




FIG. 238 10a




FIG. 235 10b

19/45




FIG. 23c 10c

20/45




FIG. 24a 11a




FIG. 24b 11b

21/45




FIG. ~~21c~~ 11c

22/45




FIG. ~~22a~~ 12a




FIG. ~~25b~~ 12b

23/45




FIG. 25c 12c




FIG. 25d 12d

24/45




FIG. 26 13




FIG. 27 14

25/45




FIG. ~~28a~~ 15a




FIG. ~~28b~~ 15b

26/45




FIG. ~~28c~~ 15c

27/45




Figure 16 (Prior Art)

28/45




Figure 17 (Prior Art)

29/45



Figure 18

30/45



Figure 1a (Prior Art)




Figure 20a (Prior Art)




Figure 20 b (Prion Ant)




Figure 20c (Prior Art)




Figure 21a




Figure 21b




Figure 21c




Figure 22a




Figure 22b



Figure 23

40/45




Figure 24

41/45




Figure 25

42/45




Figure 26

43/45




Figure 27



Figure 28




Figure 29

INTERNATIONAL SEARCH REPORT

International Application No

PCT/CA 97/00907

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06T5/00

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 G06T

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

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 454 371 A (FENSTER AARON ET AL) 3 October 1995 see abstract see column 6, line 12 - line 35 see column 8, line 44 - column 10, line 14	1-5
Y	see column 14, line 1 - column 15, line 29	6-12
Y	EP 0 612 502 A (S S B CO LTD) 31 August 1994 see abstract see column 2, line 1 - line 25 see column 3, line 42 - line 45 see column 8, line 4 - line 9; figure 6 see column 14, line 53 - column 25, line 13	6-12
	--- -/--	

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search

17 April 1998

Date of mailing of the international search report

24/04/1998

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Gonzalez Ordonez, O

INTERNATIONAL SEARCH REPORT

International Application No

PCT/CA 97/00907

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
P,A	US 5 682 895 A (ISHIGURO MASAOKI) 4 November 1997 see abstract see column 2, line 16 - line 35 see column 11, line 1 - column 14, line 7 -----	1-12

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/CA 97/00907

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5454371 A	03-10-95	AU 2730495 A	19-01-96
		CA 2193485 A	04-01-96
		WO 9600425 A	04-01-96
		EP 0766857 A	09-04-97
		JP 10502194 T	24-02-98
		WO 9600402 A	04-01-96
		GB 2308661 A	02-07-97
		US 5562095 A	08-10-96
EP 0612502 A	31-08-94	JP 6245937 A	06-09-94
		AU 656288 B	27-01-95
		AU 5524094 A	22-09-94
		US 5413106 A	09-05-95
US 5682895 A	04-11-97	JP 9103432 A	22-04-97