the directly adjacent neighboring voxels v has the lowest value for $1-|n_u \cdot n_v|$, where n_u is the normal vector for the voxel u and n_v is the normal vector for a neighboring voxel v (process action **802**). It is then ascertained if the identified neighboring voxel has a previously determined sign for its plane normal as indicated by process action **804**. If so, the same sign is applied to the plane normal of the "undetermined" voxel u (process action **806**). However, if the neighboring voxel having the lowest value for $1-|n_u \cdot n_v|$ also has a plane normal with an undetermined sign, then the identified neighboring voxel is designated as the currently selected voxel (process action **810**) as long as a prescribed propagation limit has not been reach (process action **808**). Process actions **802** through **810** are repeated until a voxel with a determined normal direction is encountered or it is determined that the prescribed propagation limit has been reached. In the case where a voxel having a determined normal sign is reached within the propagation limits, the sign associated with that voxel's normal is assigned to the currently selected voxel as well as any previously selected "undetermined" voxels that were traversed on the way to the currently selected voxel (process action **806**). The aforementioned prescribed propagation limit (or so-called marching distance) is imposed to prevent excessive propagation to parts of the object surface which may not be related to the surface in the ambiguous region. This limit will vary depending on the object being modeled. If in process action **808**, it is ascertained that the prescribed propagation limit has been reached, then the "undetermined sign" status of the currently selected voxel, and that of any previously selected "undetermined" voxels that were traversed on the way to the currently selected voxel, are retained (process action **812**). The process of action **800** through **806** or **812** is then repeated for any remaining previously unselected voxels whose normal vector was assigned an "undetermined" sign as indicated by process action **814**.

Once the foregoing back propagation procedure is complete, it is preferred that an additional check be made to ensure the derived normal signs are realistic and to designate a sign for the normal vectors of any remaining "undetermined" voxels. Specifically, this entails imposing a local consistency check. Referring to FIG. **9**, the check procedure begins by selecting one of the voxels (process action **900**). A prescribed number of voxels neighboring the selected voxel are then identified in process action **902**. As before, a fixed block (e.g., a 3 by 3 by 3 voxel block) can be employed, but the size of the neighborhood could be chosen via any appropriate procedure. The process continues with the identification of the signs associated with the normals, if previously determined, of the selected voxel and its neighboring voxels (process action **904**). The sign associated with the majority of the voxels is then determined (process action **906**). The sign of the normal vector of the selected voxel is re-assigned to match the majority sign if it is opposed, or is assigned the majority sign if the normal vector's sign was previously designated as undetermined (process action **908**). Process action **900** through **908** are repeated for each remaining unselected voxel, as indicated by process action **910**. The local consistency check ends the aforementioned signed distance computation portion of the preferred surface extraction method. At this point, the signed direction and magnitude of the plane normals for each voxel have been computed.

Referring once again to FIG. **5**, the next action in the surface extraction procedure is to extract a triangular-mesh representation of the surface of the object defined by the previously computed planes within each voxel associated with the object's surface (process action **506**). Preferably, a modified marching cubes approach is employed for this purpose. The marching cubes procedure is a known method for extracting a triangular-mesh representation from an implicit surface. This method is modified in the present object modeling process to incorporate an octree based approach, which has the advantage of not needing a seed cube and eliminating the possibility that a disjointed portion of the object being modeled would be missed. In the traditional marching cubes method, a voxel is selected (i.e., the seed cube) that contains a portion of the surface of the object. The triangular-mesh representation of the surface is then constructed by, inter alia, identifying a neighboring voxel containing a part of the surface and "marching" voxel by voxel until surface is analyzed. However, if the object being modeled is made up of multiple, separated sections, then the traditional method will miss the disjointed portions of the object not containing the surface component of the seed cube. In the modified procedure, this problem is resolved using the aforementioned octree approach. As explained earlier, the point cloud associated with the registered 3D reconstructions was partitioned into voxels using the hierarchical octree approach that ultimately identified each voxel that contained reconstruction points. The modified marching cubes procedure uses this information as follows. Referring to FIG. **10**, any one of the previously defined voxels containing reconstruction points is selected (process action **1000**). The triangular-mesh surface representation is then computed by proceeding as in the traditional marching cubes method, with the exception that an accounting of each voxel processed is kept (process action **1002**). When all the voxels containing a section of the surface of the portion of the object associated with the initially selected voxel have been processed, it is determined in process action **1004** if any un-processed voxels exist (as would be the case if the object is made up of disjointed surfaces). Whenever, it is determined that un-processed voxels exist, one of them is selected (process action **1006**) and the procedure of process actions **1002** through **1006** is repeated until no more un-processed voxels remain. In this way every voxel containing reconstruction points is processed and so every surface making up the object will be modeled.

5.0 Texture Mapping.

Once the surface extraction procedure is complete, it is possible to also perform a texture mapping process to create a photorealistic model of the image. While many texture mapping method exist that could be used for this purpose, it is preferred that texture be added to the model using a method of blending portions of the original images used to produce the 3D reconstructions as was described in a co-pending patent application entitled PATCH-BASED ALIGNMENT METHOD AND APPARATUS FOR CONSTRUCTION OF IMAGE MOSAICS and having a common assignee with the present application. This co-pending application was filed on Aug. 8, 1997, and assigned Ser. No. 08/905,100. The disclosure of the co-pending application is hereby incorporated by reference. Essentially, for each of the triangular areas of the object's surface defined in the surface extraction procedure, the portions of the original images that depict that particular triangular area are identified via known methods. These areas are then blended to create a composited representation of the area. Finally, this composited representation is assigned as the texture for the selected area.

6.0 Experimental Results.

In a tested embodiment of the present invention, a commercial trinocular stereo system was used to capture images