where stereo matching techniques have been used to generate the reconstructions as they tend to be noisy. One preferred noise elimination procedure is "automatic clustering". This method begins by calculating the mean and then the variance of the points in the point cloud in each orthogonal direction. A threshold is then applied in each direction to define a cuboid-shaped bounding box. Any point lying outside the bounding box is eliminated.

The procedure is repeated until there is no longer any significant change in the mean and variance between iterations. While the automatic clustering method is good at eliminating extraneous reconstruction points outside the defined bounding boxes, it does nothing to eliminate extraneous points within the box, such as might exist in voids or holes associated with the object. A "3D spatial filtering" procedure can be used to remove such points from the reconstruction data. The 3D spatial filtering begins by dividing a 3D space containing all the reconstruction points into voxels. To minimize processing, an octree scheme [**3**] is employed resulting in only voxels containing reconstruction points being considered. For each point, the voxel containing the point is identified along with a prescribed number of its neighboring voxels. All the points contained in the voxel block are counted, and if the total number of points exceeds a prescribed threshold, then the point remains in the reconstruction data. Otherwise, the point is eliminated.

Once the noise elimination processing is complete, the various individual 3D reconstructions are merged into one overall 3D reconstruction of the object via a standard registration process. This registration is required because the images used to compute the individual 3D reconstructions would have been captured at a different orientations in relation to the object. Thus, the coordinate frames of each group may be different, and so to create an overall 3D reconstruction, the point sets associated with each of the individual reconstructions have to be aligned to a common coordinate frame.

The next phase of the object modeling process involves a surface extraction procedure designed to define the surface of the object based on the points associated with the previously-computed overall 3D reconstruction. One preferred procedure for accomplishing this task begins by dividing a 3D space containing all the reconstruction points associated with the overall 3D reconstruction into voxels using an octree approach so that only those voxels containing at least on reconstruction point are identified. Each voxel in turn undergoes a "signed distance computation" to define a plane which best represents the surface of the object in that voxel. Specifically, the signed distance computation begins by identifying a "neighborhood" of voxels associated with the voxel under consideration. For the purposes of the present object modeling process, a fixed neighborhood size was used—for example, a 3 by 3 by 3 voxel block as used in tested embodiments of the present invention. All the points contained within the identified voxels neighborhood are used to calculate a plane that represents the surface of the object contained within the voxel under consideration. This procedure is then repeated for all the voxel containing reconstruction points.

Preferably, the plane for each voxel is defined by a normal thereof extending from the plane to a prescribed one of the vertices of the voxel under consideration. This normal is preferably established by first computing the centroid of all points in the previously identified voxel neighborhood. A covariance matrix is then computed and the eigenvector corresponding to the smallest eigenvalue of the covariance matrix is designated as vector of the normal of the plane, but without initially specifying which of the two possible directions that the vector is directed. The distance from the plane to the prescribed vertex along the normal is also calculated to establish the magnitude of the normal vector. The direction of the normal for each voxel is preferably established by first identifying the direction from each point in the voxel under consideration to the optical center associated with the camera used to capture the original image from which the point was derived. This is repeated for each of the other voxels in the previously identified voxel neighborhood associated with the voxel under consideration. The vector from a reconstruction point to its associated optical center is referred to as the visibility vector. It is next determined whether the angle between the normal computed for each voxel in a voxel neighborhood and the visibility vector for each point contained in that voxel is less than 90 degrees by a prescribed threshold amount, greater than 90 degrees by the prescribed threshold amount, or within the prescribed threshold amount of 90 degrees. The normal vector of a voxel under consideration would be assigned a positive direction (i.e., toward the prescribed vertex) if a majority of the angles between the visibility vectors and associated normals were less than 90 degrees by the threshold amount, and assigned a negative direction (i.e., away from the prescribed vertex) when the majority of the angles are greater than 90 degrees by the threshold amount. However, there can be regions of an object's surface where the plane normal of a voxel is almost perpendicular to most of the visibility vectors associated with the points contained in the voxel. If this occurs there is some ambiguity as to whether the normal points in the positive or negative direction. Thus, if the majority of the of the angles between the visibility vectors and associated normals were found to be within the prescribed threshold amount of 90 degrees, then the undetermined direction status is maintained for the normal of the voxel under consideration.

Once the foregoing procedure has been performed for all the voxels containing reconstruction points, the direction of the normal of any voxels still having an undetermined normal direction status is preferably determined by a back propagation procedure. The back propagation procedure begins by selecting a voxel that was marked as having a plane normal with an undetermined direction, and identifying which of the directly adjacent neighboring voxels has the largest absolute value for the cross product of the normal vector associated with the currently selected voxel and the normal vector associated with the neighboring voxel. If the so identified neighboring voxel has a previously determined direction for its plane normal, then the same direction is assigned to the plane normal of the "undetermined" voxel. However, if the identified neighboring voxel also has a plane normal with an undetermined direction, then the identified neighboring voxel becomes the currently selected voxel and the process is repeated until a voxel with a determined normal direction is encountered or a prescribed propagation limit is reached. In the case where a voxel having a determined normal direction is reached within the propagation limits, the direction associated with that voxel's normal is assigned to all the "undetermined" voxels traversed on the way to the "determined" voxel. If the prescribed propagation limit is reached before encountering a voxel having a determined direction, then the undetermined normal direction status of the currently selected voxel, and that of any "undetermined" voxels that were traversed on the way to the currently selected voxel, are retained.

Once the back propagation procedure is complete, it is preferred that an additional check be made to ensure the