

AUTOMATIC 3D MODELING SYSTEM AND METHOD

Related Application

This application claims priority under 35 USC § 119 from U.S. Provisional Patent Application Serial No. 60/312,384 filed on August 14, 2001 and entitled "Automatic 3D Modeling System And Method" which is incorporated herein by reference.

5 Field of the Invention

The present invention is related to 3D modeling systems and methods and more particularly, to a system and method that merges automatic image-based model generation techniques with interactive real-time character orientation techniques to provide rapid creation of virtual 3D personalities.

10 Background Of The Invention

There are many different techniques for generating an animation of a three dimensional object on a computer display. Originally, the animated figures (for example, the faces) looked very much like wooden characters, since the animation was not very good. In particular, the user would typically see an animated face, yet its features and expressions would be static. Perhaps
15 the mouth would open and close, and the eyes might blink, but the facial expressions, and the animation in general, resembled a wood puppet. The problem was that these animations were typically created from scratch as drawings, and were not rendered using an underlying 3D model to capture a more realistic appearance, so that the animation looked unrealistic and not very life-like. More recently, the animations have improved so that a skin may cover the bones of the
20 figure to provide a more realistic animated figure.

While such animations are now rendered over one or more deformation grids to capture a more realistic appearance for the animation, often the animations are still rendered by professional companies and redistributed to users. While this results in high-quality animations, it is limited in that the user does not have the capability to customize a particular animation, for
25 example, of him or herself, for use as a virtual personality. With the advance features of the Internet or the World Wide Web, these virtual personas will extend the capabilities and

interaction between users. It would thus be desirable to provide a 3D modeling system and method which enables the typical user to rapidly and easily create a 3D model from an image, such as a photograph, that is useful as a virtual personality.

Typical systems also required that, once a model was created by the skilled animator, the same animator was required to animate the various gestures that you might want to provide for the model. For example, the animator would create the animation of a smile, a hand wave or speaking which would then be incorporated into the model to provide the model with the desired gestures. The process to generate the behavior/gesture data is slow and expensive and requires a skilled animator. It is desirable to provide an automatic mechanism for generating gestures and behaviors for models without the assistance of a skilled animator. It is to these ends that the present invention is directed.

Summary Of The Invention

Broadly, the invention utilizes image processing techniques, statistical analysis and 3D geometry deformation to allow photo-realistic 3D models of objects, such as the human face, to be automatically generated from an image (or from multiple images). For example, for the human face, facial proportions and feature details from a photograph (or series of photographs) are identified and used to generate an appropriate 3D model. Image processing and texture mapping techniques also optimize how the photograph(s) is used as detailed, photo-realistic texture for the 3D model.

In accordance with another aspect of the invention, a gesture of the person may be captured and abstracted so that it can be applied to any other model. For example, the animated smile of a particular person may be captured. The smile may then be converted into feature space to provide an abstraction of the gesture. The abstraction of the gesture (e.g., the movements of the different portions of the model) are captured as a gesture. The gesture may then be used for any other model. Thus, in accordance with the invention, the system permits the generation of a gesture model that may be used with other models.

In accordance with the invention, a method for generating a three dimensional model of an object from an image is provided. The method comprises determining the boundary of the object to be modeled and determining the location of one or more landmarks on the object to be

modeled. The method further comprises determining the scale and orientation of the object in the image based on the location of the landmarks, aligning the image of the object with the landmarks with a deformation grid, and generating a 3D model of the object based on the mapping of the image of the object to the deformation grid.

5 In accordance with another aspect of the invention, a computer implemented system for generating a three dimension model of an image is provided. The system comprises a three dimensional model generation module further comprising instructions that receive an image of an object and instructions that automatically generate a three dimensional model of the object. The system further comprises a gesture generation module further comprising instructions for
10 generating a feature space and instructions for generating a gesture object corresponding to a gesture of the object so that the gesture behavior may be applied to another model of an object.

In accordance with yet another aspect of the invention, a method for automatically generating an automatic gesture model is provided. The method comprises receiving an image of an object performing a particular gesture and determining the movements associated with the
15 gesture from the movement of the object to generate a gesture object wherein the gesture object further comprises a coloration change variable storing the change of coloration that occur during the gesture, a two dimensional change variable storing the change of the surface that occur during the gesture and a three dimensional change variable storing the change of the vertices associated with the object during the gesture.

20 In accordance with yet another aspect of the invention, a gesture object data structure that stores data associated with a gesture for an object is provided. The gesture object comprises a texture change variable storing changes in coloration of a model during a gesture, a texture map change variable storing changes in the surface of the model during the gesture, and a vertices change variable storing changes in the vertices of the model during the gesture wherein the
25 texture change variable, the texture map change variable and the vertices change variable permit the gesture to be applied to another model having a texture and vertices. The gesture object data structure stored its data in a vector space where coloration, surface motion and 3D motion may be used by many individual instances of the model.

Brief Description Of The Drawings

30 Figure 1 is a flowchart describing a method for generating a 3D model of a human face;

Figure 2 is a diagram illustrating an example of a computer system which may be used to implement the 3D modeling method in accordance with the invention;

Figure 3 is a block diagram illustrating more details of the 3D model generation system in accordance with the invention;

5 Figure 4 is an exemplary image of a person's head that may be loaded into the memory of a computer during an image acquisition process;

Figure 5 illustrates the exemplary image of Figure 4 with an opaque background after having processed the image with a "seed fill" operation;

10 Figure 6 illustrates the exemplary image of Figure 5 having dashed lines indicating particular bound areas about the locations of the eyes;

Figure 7 illustrates the exemplary image of Figure 6 with the high contrast luminance portion of the eyes identified by dashed lines;

Figure 8 is an exemplary diagram illustrating various landmark location points for a human head;

15 Figure 9 illustrates an example of a human face 3D model in accordance with the invention;

Figures 10A-10D illustrate respective deformation grids that can be used to generate a 3D model of a human head;

Figure 10E illustrates the deformation grids overlaid upon one another;

20 Figure 11 is a flowchart illustrating the automatic gesture behavior generation method in accordance with the invention;

Figure 12A and 12B illustrate an exemplary psuedo-code for performing the image processing techniques of the invention;

25 Figures 13A and 13B illustrate an exemplary work flow process for automatically generating a 3D model in accordance with the invention;

Figures 14A and 14B illustrate an exemplary pseudo-code for performing the automatic gesture behavior model in accordance with the invention;

Figure 15 illustrates an example of a base 3D model for a first model, Kristen;

Figure 16 illustrates an example of a base 3D model for a second model, Ellie;

5 Figure 17 is an example of the first model in a neutral gesture;

Figure 18 is an example of the first model in a smile gesture;

Figure 19 is an example of a smile gesture map generated from the neutral gesture and the smile gesture of the first model;

10 Figure 20 is an example of the feature space with both the models overlaid over each other;

Figure 21 is an example of a neutral gesture for the second model; and

Figure 22 is an example of the smile gesture, generated from the first model, being applied to the second model to generate a smile gesture in the second model.

Detailed Description Of The Preferred Embodiment

15 While the invention has a greater utility, it will be described in the context of generating a 3D model of the human face and gestures associated with the human face. Those skilled in the art recognize that any other 3D models and gestures can be generated using the principles and techniques described herein, and that the following is merely exemplary of a particular application of the invention and the invention is not limited to the facial models described herein.

20 To generate a 3D model of the human face, the invention preferably performs a series of complex image processing techniques to determine a set of landmark points 10 which serve as guides for generating the 3D model. Figure 1 is a flow chart describing a preferred algorithm for generating a 3D model of a human face. With reference to Figure 1, an image acquisition process (Step 1) is used to load a photograph(s) (or other image) of a human face (for example, a
25 "head shot") into the memory of a computer. Preferably, images may be loaded as JPEG images, however, other image type formats may be used without departing from the invention. Images

can be loaded from a diskette, downloaded from the Internet, or otherwise loaded into memory using known techniques so that the image processing techniques of the invention can be performed on the image in order to generate a 3D model.

Since different images may have different orientations, the proper orientation of the image should be determined by locating and grading appropriate landmark points 10. Determining the image orientation allows a more realistic rendering of the image onto the deformation grids. Locating the appropriate landmark points 10 will now be described in detail.

Referring to Figure 1, to locate landmark points 10 on an image, a “seed fill” operation may preferably be performed (Step 2) on the image to eliminate the variable background of the image so that the boundary of the head (in the case of a face) can be isolated on the image. Figure 4 is an exemplary image 20 of a person’s head that may be loaded into the memory of the computer during the image acquisition process (Step 1, Figure 1). A “seed fill” operation (Step 2, Figure 1) is a well-known recursive paintfill operation that is accomplished by identifying one or more points 22 in the background 24 of the image 20 based on, for example, color and luminosity of the point(s) 22 and expand a paintfill zone 26 outwardly from the point(s) 22 where the color and luminosity are similar. Preferably, the “seed fill” operation successfully replaces the color and luminescent background 24 of the image with an opaque background so that, the boundary of the head can be more easily determined.

Referring again to Figure 1, the boundary of the head 30 can be determined (Step 3), for example, by locating the vertical center of the image (line 32) and integrating across a horizontal area 34 from the centerline 32 (using a non-fill operation) to determine the width of the head 30, and by locating the horizontal center of the image (line 36) and integrating across a vertical area 38 from the centerline 36 (using a non-fill operation) to determine the height of the head 30. In other words, statistically directed linear integration of a field of pixels whose values differ based on the presence of an object or the presence of a background is performed. This is shown in Figure 5 which shows the exemplary image 20 of Figure 4 with an opaque background 24.

Returning again to Figure 1, upon determining the width and height of the head 30, the bounds of the head 30 can be determined by using statistical properties of the height of the head 30 and the known properties of the integrated horizontal area 34 and top of the head 30. Typically, the height of the head will be approximately 2/3 of the image height and the width of

the head will be approximately 1/3 of the image width. The height of the head may also be 1.5 times the width of the head which is used as a first approximation.

Once the bounds of the head 30 are determined, the location of the eyes 40 can be determined (Step 4). Since the eyes 40 are typically located on the upper half of the head 30, a statistical calculation can be used and the head bounds can be divided into an upper half 42 and a lower half 44 to isolate the eye bound areas 46a, 46b. The upper half of the head bounds 42 can be further divided into right and left portions 46a, 46b to isolate the left and right eyes 40a, 40b, respectively. This is shown in detail in Figure 6 which shows the exemplary image 20 of Figure 4 with dashed lines indicating the particular bound areas.

Referring yet again to Figure 1, the centermost region of each eye 40a, 40b can be located (Step 5) by identifying a circular region 48 of high contrast luminance within the respective eye bounds 46a, 46b. This operation can be recursively performed outwardly from the centermost point 48 over the bounded area 46a, 46b and the results can be graded to determine the proper bounds of the eyes 40a, 40b. Figure 7 shows the exemplary image of Figure 6 with the high contrast luminance portion of the eyes identified by dashed lines.

Referring again to Figure 1, once the eyes 40a, 40b have been identified, the scale and orientation of the head 30 can be determined (Step 6) by analyzing a line 50 connecting the eyes 40a, 40b to determine the angular offset of the line 50 from a horizontal axis of the screen. The scale of the head 30 can be derived from the width of the bounds according to the following formula: width of bound/width of model.

After determining the above information, the approximate landmark points 10 on the head 30 can be properly identified. Preferred landmark points 10 include a) outer head bounds 60a, 60b, 60c; b) inner head bounds 62a, 62b, 62c, 62d; c) right and left eye bounds 64a-d, 64w-z, respectively; d) corners of nose 66a, 66b; and e) corners of mouth 68a, 68b (mouth line), however, those skilled in the art recognize that other landmark points may be used without departing from the invention. Figure 8 is an exemplary representation of the above landmark points shown for the image of Figure 4.

Having determined the appropriate landmark locations 10 on the head 30, the image can be properly aligned with one or more deformation grids (described below) that define the 3D

model 70 of the head (Step 7). The following describes some of the deformation grids that may be used to define the 3D model 70, however, those skilled in the art recognize that these are merely exemplary of certain deformation grids that may be used to define the 3D model and that other deformation grids may be used without departing from the invention. Figure 9 illustrates an example of a 3D model of a human face generated using the 3D model generation method in accordance with the invention. Now, more details of the 3D model generation system will be described.

Figure 2 illustrates an example of a computer system 70 in which the 3D model generation method and gesture model generation method may be implemented. In particular, the 3D model generation method and gesture model generation method may be implemented as one or more pieces of software code (or compiled software code) which are executed by a computer system. The methods in accordance with the invention may also be implemented on a hardware device in which the method in accordance with the invention are programmed into a hardware device. Returning to Figure 2, the computer system 70 shown is a personal computer system. The invention, however, may be implemented on a variety of different computer systems, such as client/server systems, server systems, workstations, etc... and the invention is not limited to implementation on any particular computer system. The illustrated computer system may include a display device 72, such as a cathode ray tube or LCD, a chassis 74 and one or more input/output devices, such as a keyboard 76 and a mouse 78 as shown, which permit the user to interact with the computer system. For example, the user may enter data or commands into the computer system using the keyboard or mouse and may receive output data from the computer system using the display device (visual data) or a printer (not shown), etc. The chassis 74 may house the computing resources of the computer system and may include one or more central processing units (CPU) 80 which control the operation of the computer system as is well known, a persistent storage device 82, such as a hard disk drive, an optical disk drive, a tape drive and the like, that stores the data and instructions executed by the CPU even when the computer system is not supplied with power and a memory 84, such as DRAM, which temporarily stores data and instructions currently being executed by the CPU and loses its data when the computer system is not being powered as is well known. To implement the 3D model generation and gesture generation methods in accordance with the invention, the memory may store a 3D modeler 86 which is a series of instructions and data being executed by the CPU 80 to implement

the 3D model and gesture generation methods described above. Now, more details of the 3D modeler will be described.

Figure 3 is a diagram illustrating more details of the 3D modeler 86 shown in Figure 2. In particular, the 3D modeler includes a 3D model generation module 88 and a gesture generator module 90 which are each implemented using one or more computer program instructions. The pseudo-code that may be used to implement each of these modules is shown in Figures 12A – 12B and Figures 14A and 14B. As shown in Figure 3, an image of an object, such as a human face is input into the system as shown. The image is fed into the 3D model generation module as well as the gesture generation module as shown. The output from the 3D model generation module is a 3D model of the image which has been automatically generated as described above. The output from the gesture generation module is one or more gesture models which may then be applied to and used for any 3D model including any model generate by the 3D model generation module. The gesture generator is described in more detail below with reference to Figure 11. In this manner, the system permits 3D models of any object to be rapidly generated and implemented. Furthermore, the gesture generator permits one or more gesture models, such as a smile gesture, a hand wave, etc...) to be automatically generated from a particular image. The advantage of the gesture generator is that the gesture models may then be applied to any 3D model. The gesture generator also eliminates the need for a skilled animator to implement a gesture. Now, the deformation grids for the 3D model generation will be described.

Figures 10A-10D illustrate exemplary deformation grids that may be used to define a 3D model 70 of a human head. Figure 10A illustrates a bounds space deformation grid 72 which is preferably the innermost deformation grid. Overlaying the bounds space deformation grid 72 is a feature space deformation grid 74 (shown in Figure 10B). An edge space deformation grid 76 (shown in Figure 10C) preferably overlays the feature space deformation grid 74. Figure 10D illustrates a detail deformation grid 7D that is preferably the outermost deformation grid.

The grids are preferably aligned in accordance with the landmark locations 10 (shown in Figure 10E) such that the head image 30 will be appropriately aligned with the deformation grids when its landmark locations 10 are aligned with the landmark locations 10 of the deformation grids. To properly align the head image 30 with the deformation grids, a user may manually refine the landmark location precision on the head image (Step 8), for example by using the

mouse or other input device to “drag” a particular landmark to a different area on the image 30. Using the new landmark location information, the image 30 may be modified with respect to the deformation grids as appropriate (Step 9) in order to properly align the head image 30 with the deformation grids. A new model state can then be calculated, the detail grid 78 can then be
5 detached (Step 10), behaviors can be scaled for the resulting 3D model (Step 11), and the model can be saved (Step 12) for use as a virtual personality. Now, the automatic gesture generation in accordance with the invention will be described in more detail.

Figure 11 is a flowchart illustrating an automatic gesture generation method 100 in accordance with the invention. In general, the automatic gesture generation results in a gesture
10 object which may then be applied to any 3D model so that a gesture behavior may be rapidly generated and reused with other models. Usually, there may need to be a separate gesture model for different types of 3D models. For example, a smile gesture may need to be automatically generated for a human male, a human female, a human male child and a human female child in order to make the gesture more realistic. The method begins is step 102 in which a common
15 feature space is generated. The feature space is common space that is used to store and represent an object image, such as a face, movements of the object during a gesture and object scalars which capture the differences between different objects. The gesture object to be generated using this method also stores a scalar field variable that stores the mapping between a model space and the feature space that permits transformation of motion and geometry data. The
20 automatic gesture generation method involves using a particular image of an object, such as a face, to generate an abstraction of a gesture of the object, such as a smile, which is then stored as a gesture object so that the gesture object may then be applied to any 3D model.

Returning to Figure 11, in step 104, the method determines the correlation between the feature space and the image space to determine the texture map changes which represent changes
25 to the surface movements of the image during the gesture. In step 106, the method updates the texture map from the image (to check the correlation) and applies the resultant texture map to the feature space and generates a variable “stDeltaChange” as shown in the exemplary pseudo-code shown in Figures 14A and 14B which stores the texture map changes. In step 108, the method determines the changes in the 3D vertices of the image model during the gesture which captures
30 the 3D movement that occurs during the gesture. In step 110, the vertex changes are applied to the feature space and are captured in the gesture object in a variable “VertDeltaChange” as

shown in Figure 14A and 14B. In step 112, the method determines the texture coloration that occurs during the gesture and applies it to the feature space. The texture coloration is captured in the "DeltaMap" variable in the gesture object. In step 114, the gesture object is generated that includes the "stDeltaChange", "VertDeltaChange" and "DeltaMap" variables which contain the coloration, 2D and 3D movement that occurs during the gesture. The variables represent only the movement and color changes that occurs during a gesture so that the gesture object may then be applied to any 3D model. In essence, the gesture object distills the gesture that exists in a particular image model into an abstract object that contains the essential elements of the gesture so that the gesture may then be applied to any 3D model.

The gesture object also includes a scalar field variable storing the mapping between a feature space of the gesture and a model space of a model to permit transformation of the geometry and motion data. The scalerArray has an entry for each geometry vertex in the Gesture object. Each entry is a 3 dimensional vector that holds the change in scale for that vertex of the Feature level from its undeformed state to the deformed state. The scale is computed by vertex in Feature space by evaluating the scaler change in distance from that vert to connected verticies. The scaler for a given Gesture vertex is computed by weighted interpolation of that Vertex's position when mapped to UV space of a polygon in the Feature Level. The shape and size of polygons in the feature level are chosen to match areas of similarly scaled movement. This was determined by analyzing visual flow of typical facial gestures. The above method is shown in greater detail in the pseudo-code shown in Figure 14A and 14B.

Figures 12A- B and Figures 13A and B, respectively, contain a sample pseudo code algorithm and exemplary work flow process for automatically generating a 3D model in accordance with the invention.

The automatically generated model can incorporate built-in behavior animation and interactivity. For example, for the human face, such expressions include gestures, mouth positions for lip syncing (visemes), and head movements. Such behaviors can be integrated with technologies such as automatic lip syncing, text-to-speech, natural language processing, and speech recognition and can trigger or be triggered by user or data driven events. For example, real-time lip syncing of automatically generated models may be associated with audio tracks. In addition, real-time analysis of the audio spoken by an intelligent agent can be provided and

synchronized head and facial gestures initiated to provide automatic, lifelike movements to accompany speech delivery.

Thus, virtual personas can be deployed to serve as an intelligent agent that may be used as an interactive, responsive front-end to information contained within knowledge bases, customer resource management systems, and learning management systems, as well as entertainment applications and communications via chat, instant messaging, and e-mail. Now, examples of a gesture being generated from an image of a 3D model and then applied to another model in accordance with the invention will now described.

Figure 15 illustrates an example of a base 3D model for a first model, Kristen. The 3D model shown in Figure 15 has been previously generated as described above using the 3D model generation process. Figure 16 illustrates a second 3D model generated as described above. These two models will be used to illustrate the automatic generation of a smile gesture from an existing model to generate a gesture object and then the application of that generated gesture object to another 3D model. Figure 17 show an example of the first model in a neutral gesture while Figure 18 shows an example of the first model in a smile gesture. The smile gesture of the first model is then captured as described above. Figure 19 illustrates an example of the smile gesture map (the graphical version of the gesture object described above) that is generated from the first model based on the neutral gesture and the smile gesture. As described above, the gesture map abstracts the gesture behavior of the first model into a series of coloration changes, texture map changes and 3D vertices changes which can then be applied to any other 3D model that has texture maps and 3D vertices. Then, using this gesture map (which includes the variables described above), the gesture object may be applied to another model in accordance with the invention. In this manner, the automatic gesture generation process permits various gestures for a 3D model to be abstracted and then applied to other 3D models.

Figure 20 is an example of the feature space with both the models overlaid over each other to illustrate that the feature space of the first and second models are consistent with each other. Now, the application of the gesture map (and therefore the gesture object) to another model will be described in more detail. In particular, Figure 21 illustrates the neutral gesture of the second model. Figure 22 illustrates the smile gesture (from the gesture map generated by

from the first model) applied to the second model to provide a smile gesture to that second model even when the second model does not actually show a smile.

While the above has been described with reference to a particular method for locating the landmark location points on the image and a particular method for generating gestures, those skilled in the art will recognize that other techniques may be used without departing from the invention as defined by the appended claims. For example, techniques such as pyramid transforms which uses a frequency analysis of the image by down sampling each level and analyzing the frequency differences at each level can be used. Additionally, other techniques such as side sampling and image pyramid techniques can be used to process the image.

Furthermore, quadrature (low pass) filtering techniques may be used to increase the signal strength of the facial features, and fuzzy logic techniques may be used to identify the overall location of a face. The location of the landmarks may then be determined by a known corner finding algorithm.

CLAIMS:

1. A method for generating a three dimensional model of an object from an image, the method comprising:

determining the boundary of the object to be modeled;

5 determining the location of one or more landmarks on the object to be modeled;

determining the scale and orientation of the object in the image based on the location of the landmarks;

aligning the image of the object with the landmarks with a deformation grid; and

generating a 3D model of the object based on the mapping of the image of the object to the deformation grid.

2. The method of Claim 1, wherein the boundary determining further comprises statistically directed linear integration of a field of pixels whose values differ based on the presence of an object or the presence of a background.

3. The method of Claim 1, wherein the boundary determining further comprises performing a statistically directed seed fill operation in order to remove the background around the image of the object.

4. The method of Claim 3, wherein determining the landmarks further comprises identifying features found by procedural correlation or band pass filtering and thresholding in statistically characteristic zones as determined during the boundary determination.

5. The method of Claim 4, wherein determining the landmarks further comprises determining additional landmarks based on a refinement of the boundary areas.

6. The method of Claim 5, wherein determining the landmarks further comprises adjusting the landmarks by the user.

7. A computer implemented system for generating a three dimension model of an image, the computer implemented system comprising:

a three dimensional model generation module further comprising instructions that receive an image of an object and instructions that automatically generate a three dimensional model of the object; and

a gesture generation module further comprising instructions for generating a feature space and instructions for generating a gesture object corresponding to a gesture of the object so that the gesture behavior may be applied to another model of an object.

8. A method for automatically generating an automatic gesture model, the method comprising:

receiving an image of an object performing a particular gesture; and

determining the movements associated with the gesture from the movement of the object to generate a gesture object wherein the gesture object further comprises a coloration change variable storing the change of coloration that occur during the gesture, a two dimensional change variable storing the change of the surface that occur during the gesture and a three dimensional change variable storing the change of the vertices associated with the object during the gesture.

9. The method of Claim 8 further comprises generating a feature space into which the gesture is mapped during the automatic gesture generation process.

10. The method of Claim 9, wherein the determining the movements further comprise determining a correlation between the feature space and the image of the object.

11. The method of Claim 9 further comprising transforming the geometric vectors and motion vectors to and from a feature space.

12. The method of Claim 9 further comprising applying changes in the coloration, texture motion and geometric motion from one model to another model using the feature space.

13. A gesture object data structure that stores data associated with a gesture for an object, comprising:

a scaler field variable storing the mapping between a feature space of the gesture and a model space of a model to permit transformation of the geometry and motion data;

a texture change variable storing changes in coloration of a model during a gesture;

a texture map change variable storing changes in the surface of the model during the gesture; and

a vertices change variable storing changes in the vertices of the model during the gesture wherein the texture change variable, the texture map change variable and the vertices change variable permit the gesture to be applied to another model having a texture and vertices.