

Extending the Range of Facial Types

Steve DiPaola

New York Institute of Technology
Wheatley Road
Computer Graphics Lab
Old Westbury, New York 11568

Abstract

We describe, in case study form, techniques to extend the range of facial types and movement using a parametric facial animation system originally developed to model and control synthetic 3D faces limited to a normal range of human shape and motion. These techniques have allowed us to create a single authoring system that can create and animate a wide range of facial types that range from realistic, stylized, cartoon-like, or a combination thereof, all from the same control system. Additionally we describe image processing and 3D deformation tools that allow for a greater range of facial types and facial animation output.

Introduction

What is the goal of a 3D facial animation system? Should its function be to create and animate faces that are scientifically accurate and therefore limited to the range of what human faces can produce or should it be able to create and control faces that are extremely cartoon-like or stylized in some way? All these types, or a combination thereof, could possibly be the correct choice depending upon the application.

A facial animation system capable of creating a limitless range of facial types would be a valuable tool for an animator. Regardless of the type of character required by a client or an application, such a system would have the capacity to animate the desired facial type, and in so doing be able to reuse animation and expression libraries from previous animation sessions, whatever the facial type formerly used. There would be no need to write one-time programs or specialized utilities for each new facial project: once such system was mastered, an artist would not have to continuously relearn animation skills on different systems.

Such a production system, as far as I know, does not yet exist. The possible realization of such tool, however, was the motivation for the creation of FAS (Facial animation system), an interactive system for modeling and animating three dimensional faces. This system, which originally was implemented on Evans and Sutherland Picture System and currently operates on a Silicon Graphics' Iris system, has over a period of time undergone a series of enhancements to broaden the range of facial types and motions available to the user.



Figure 1. This realistic head was created with FAS using standard eyebrow, eye, and mouth parameters. Warping functions were used to add details such as eye wells, accentuated nose curvature and muscle bulges under the mouth.



Figure 2. This stylized character was modeled by first creating the overall head space with the jaw and head conformation parameters, then adding the expression, and finally applying global warping and twisting functions to distort large areas of the head.

It is in the context of expanding the range of possible facial types that I will describe FAS and its successive development since its creation. Although FAS is a full animation system and animation is a fundamental concern in the creation of a facial model, for purpose of conciseness, the concentration of this case study will focus primarily on the modeling and creation of facial types rather than on various aspects of facial animation.

Facial models are created within FAS by manipulating a single input polygonal data set. This data set defines a simplified generic human head in a neutral pose that was originally digitized from a bust using photogrammetric techniques. Since all facial types within this system are derived from the same basic input set, any created facial type - realistic, stylized, cartoon-like, or a combination thereof - can interpolate to any other facial type.

PARAMETERS

FAS, based on original research by Fred Parke [1], uses a parametric approach, that is, a set of parameters act upon a connected set of vertices which form the facial model.

This animation system currently uses approximately 100 parameters to control the face, head and facial hair, which are, in most cases, applied independently to specific facial regions (i.e. mouth area, eye area, etc.). The operations of some parameters are as simple as linear transformations (move, rot, scale) manipulating a particular sub-region of the face or head. Even in such simple cases, these parameters weigh their effects by tapering the desired operation to certain vertices as they approach a region's boundary. Other parametric types use more complicated procedural or algorithmic solutions to produce the required naturalistic facial movement.




Figure 3. In the creation of this cartoon-like head, the animator chose a “hair glob” type and defined paths on the model. FAS then automatically mapped repeated ‘hair globs’ to the chosen paths with a cubic-spline technique. The facial hair flexes with the face during animation.


Whatever the parametric type chosen, the fundamental criterion in the design of the parameter set is to provide simplicity and natural facial movement for an animator. For example, if an animator manipulates the eyelid parameter, the skin of the eyelid will wrap around the eye in a natural way regardless of the other parallel operations on the eye (i.e. eyeball scale, position of eyebrow and nose). The complexities of the facial motion are solved in the parameter algorithm so that the manipulation of the face model seems natural to the animator.

The desired goal in the design of any parameterized system is to use the smallest number of parameters while simultaneously achieving the greatest amount of predictable control for the largest amount of faces and expressions - a demanding objective, but one that need not be fully realized in order to produce satisfying results. Continual refinement of the parameters, based upon varied results attained through diverse uses by animators, has become an integral factor in the development of this animation system.

Since no two artists would animate a scene identically, or for that matter, require the same tools from an animation system, the flexibility of a parameter set is, therefore, of greatest importance. This is especially true of systems that attempt to produce facial animation without being limited by physical reality.



4a



4b

Figure 4. Given the input head depicted in the center of Figure 4(a), new polygonal data sets were created using stochastic noise deformation. These data sets functioned as starting points for editing and animating new facial types in FAS.

EXTENDING PARAMETER RANGES

FAS was initially comprised of parameters which simulated realistic human movement. Most of these parameters acted upon specific facial regions. Such a system was effective so long as the value of the parameter described the normal range of human shape and motion.

What if an animator in creating character decided to scale the eye region to half the size of the head? In such a case, the eye scale parameter would also have to rescale the neighboring regions, such as the forehead and the cheek area, to compensate for such an enlarged eye. In order to accommodate the realization of such broader requirements, parameters of the elementary animation system were thus redesigned to communicate with and affect neighboring regions. In this refinement process their effective range was also extended beyond the realistic representation of human facial movement. Although this expanded development of all parameters has increased algorithmic complexity and lengthened computational time, the diversity and flexibility of unique facial types, however has also been greatly enhanced.

When designing extreme cases beyond the normal range of human shape and motion, alternative sources of facial design and movement information have been required. Animal physiology and traditional animation techniques have served as valuable reference guidelines. Squash and stretch is one such traditional animation technique that has been added to FAS. Besides greatly enhancing animation, this parameter, which preserves the volume of an object while at the same time scaling height or thickness, has been used effectively as a modeling tool for the system.

WARPING VOLUMES

Ellipsoidal warping and twisting volumes were both recently added to the parameter set, creating a new level of flexibility in this system. These warping functions are a departure from the normal type of parameter in that they do not directly affect specific vertices but are instead applied to those vertices which lie within the ellipsoidal volume of the warp.

In this technique, facial varieties are displaced by a sum of a vector-valued warp functions. Each warp is a function $f(p-0)$ of the direction vector from the warp origin 0 to the vertex p . The warp function decays smoothly to zero so that the warps have localized influence over the model.


Warping volumes can be interactively scaled and moved in three dimensions to a specific facial area or over the entire head. The animator can choose the level of warping or twisting to be applied to a particular region; multiple warps (currently a maximum of 12) are often used. Rather than adding very specific, seldom used parameters- to create a mole, pointed ears, or a double chin- warping volumes can be substituted instead. Moreover, they have strong applications for gross character design (i.e. large asymmetrical cranium, accentuated square chin), special effects (i.e. head twisted like a cork screw) as well as subtle realistic nuances (i.e. labial, nasal and brow furrows). Figures 1, 2, and 3 are examples illustrating the use of warping and standard parameters to create faces that are realistic, stylized or cartoon-like. This type of displacement function was chosen over more commonly used methods such as free-form deformation [2, 4] because of its speed for interactive work and also for its simplicity of use.

STOCHASTIC NOISE DEFORMATION

Although not an integral part of FAS, stochastic noise deformation [3] has been successfully used in conjunction with the system to perturb the vertices of an existing facial model by vector-valued solid noise. Stochastic deformation has been applied effectively in two ways. First, by slightly deforming the input data set, a large array of new input faces can be created, characterized by either realistic or caricature individuality (Figures 4 (a) and 4(b), respectively); these can then be altered and animated with FAS in normal manner. A second, more stylized approach perturbs the face once it has already been created with FAS by applying a more severe deformation onto the facial model (Figures 5(a) and 5(b)).



5a



5b

Figure 5. Once the facial models created in FAS, stochastic noise deformation was used as a post-process to deform the head models.

IMAGE PROCESSING

Image processing of 3D models is often overlooked as a technique for creating visually unique facial types. It can, however, contribute a graphic component to simple 3D facial models.

Although image processing of 3D models is conceptually analogous to video or film image processing, there are some definite advantages to applying 2D processing to each frame of a rendered 3D model, since artifacts of the 3D data can be used to create the final image. In Figure 6, the depth buffer of the facial model was composited with the image buffer of the eye and an edge enhancing filter was applied. In figure 7, different angles of the same facial model -- one smooth shaded, the other flat-shaded were

matted through a checkerboard image; color manipulation and edge enhancing were then applied to different layers.

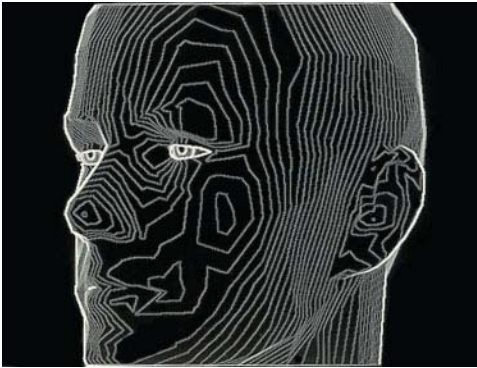


Figure 6. After the creation of the facial models, edge enhancing was applied to the depth buffer if the rendered image. This Figure is from an animated sequence of the music video “Musique Non-Stop” by Kraftwerk.

LIMITATIONS AND FUTURE IMPROVEMENTS

The topology used for the facial model is an irregular polygonal network. With this representation, polygons are sized and positioned to match the features of the face; there are many small polygons in areas of high curvature and few large ones in the flatter regions. Polygonal boundaries are designed to coincide with the desired creases in the face.

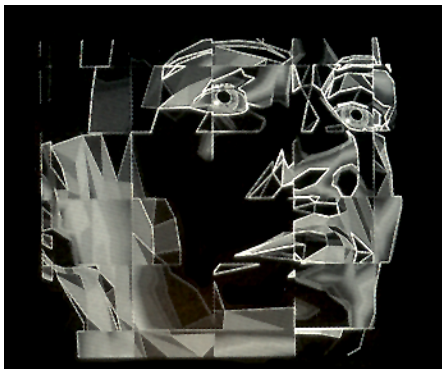


Figure 7. Different angles from the same facial model were matted through a checkerboard image. Color manipulation and edge enhancing were applied to different layers.

An irregular topology can be less dense than other topological representations while still achieving a reasonable level of surface curvature, thereby cutting down significantly on computation time; it is also capable of being more specialized than regular grid or spline topologies. However, this specialization holds the greatest limitation, since it implies explicit, direct control of specific facial vertices. Such direct control, while exhibiting certain favorable qualities, is restricting in terms of flexibility and creating generalized manipulation approaches.

By contrast, spline-based surfaces such as bi-cubic patches are capable of a more indirect approach since they require only the control points to be explicitly specified. Bi-cubic patches, however, are more difficult to refine, and additional detail must be added to the entire facial model even when it is only needed in minimal areas such as those presented by Forsey and Bartels, appear to offer what might be at the present state of animation research the best solution for obtaining indirect control of a facial model, providing variable resolution and high quality surface curvature.

Future areas of research for FAS will include the use of patches, developing a hair model that can incorporate a variety of hairstyles, and techniques to easily model and animate a large range of facial wrinkles, furrows and bulges.

References:

1. F. Parke, ‘Parameterized model for facial animation’, IEEE Computer Graphics and Applications, 2, (9), 61-88, (1982).
2. T. Sederbergh and S. Parry, ‘Freeform deformation of solid geometric models’, SIGGRAPH ’86 Conference Proceedings, 20, (2), 151-160, (1986).
3. J. Lewis, ‘Algorithms for solid noise synthesis’, SIGGRAPH ’89 Conference Proceedings, 23, (3), 263-270, (1989).
4. D. Forsey and R. Bartels, ‘Hierarchical B-spline refinement’, SIGGRAPH ’88 Conference Proceedings, 22, (4), 205-212, (1988).

Preprint Update: DiPaola Contact Info:
www.dipaola.org - steve@dipaola.org