A VIRTUAL SCULPTURE BASED MORPHABLE FACE MODEL

A Thesis

by

JESSICA LAUREN RIEWE

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

May 2007

Major Subject: Visualization Sciences

A VIRTUAL SCULPTURE BASED MORPHABLE FACE MODEL

A Thesis

by

JESSICA LAUREN RIEWE

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

Approved by:

Chair of Committee,	Frederic I. Parke
Committee Members,	Carol LaFayette
	Ricardo Gutierrez-Osuna
Head of Department,	Mark Clayton

May 2007

Major Subject: Visualization Sciences

ABSTRACT

A Virtual Sculpture Based Morphable Face Model. (May 2007)

Jessica Lauren Riewe, B.S., Texas A&M University

Chair of Advisory Committee: Dr. Frederic Parke

Exemplar virtual three-dimensional sculptures collected from sixteen artists are used to develop a multi-dimensional space to describe facial shape, which is then used as a mechanism for modeling new faces. This is accomplished through identifying and varying derived principal components of the multi-dimensional space. The relationships between these principal components and their effects on new faces are explored through the creation of new virtual faces using a graphical user interface. These new virtual face sculptures are then modified by facial feature, gender, and expression using a feature-based transformation interface based on difference vectors. Finally, animations are created to illustrate the results of these approaches. Facial mesh transformations based on principal components does not give direct predictable control for modifying specific facial features. However, it does provide interesting design choices for artists, and general patterns of the data variance were obtained. The feature-based transformations were successful in further modifying created faces.

TABLE OF CONTENTS

TABLE OF CONTENTS

LIST OF FIGURES

FIGURE Page

FIGURE Page

CHAPTER I

INTRODUCTION

Representation of human faces and their expressions is a formidable challenge because of our familiarity with faces. Yet, artists have striven to represent faces in all media over thousands of years. Interest and work in the field of computer-based three-dimensional facial representation and animation has spanned more than thirty years. In my work, exemplar virtual three-dimensional sculptures are used to develop a multi-dimensional space to describe facial shape. This space is then used as a mechanism for representing new faces.

What we recognize as a human face has a common set of features. A complete feature set includes the facial surface mask, facial feature details, and hair. A complete model will also have a back to the head, and a neck [Parke and Waters 1996]. The facial mask describes the front skin surface of the face and neck. Without details, the facial mask alone is not a very realistic representation of the face. Facial feature details include the eyes, lips, teeth, tongue, ears, and any asymmetries. Eye motion and pupil dilation are expected for successful facial animation. Lips and their resulting mouth postures are a major component of expression and speech. The tongue is required for representing and animating speech postures. Ears are very individual and needed for representing a specific person. And, while we tend to model faces with left to right symmetry, real faces will have asymmetries that occur in both conformations and expressions [Parke and Waters 1996]. Expression representation is an important aspect of faces. We recognize facial expressions that convey basic human emotion. The universal emotions are joy, sorrow, fear, surprise, disgust, and anger [Ekman 1989].

Most prior research has focused on real faces and real expressions. However, there are many applications that do not require 'real' faces. The gaming industry applies prior research to human and non-human faces alike, while animation and other art make use of caricature and virtual sculpture. Caricatures distort or exaggerated the most recognizable features of an individual's face to separate them from the average or 'normal' face.

My work presented here will focus on using artist created 3D faces, or *virtual facial sculptures*, as the basis for modeling faces. This research focuses only on the surface shape of the facial mask.

This thesis follows the style of *ACM Transactions on Graphics*.

CHAPTER II

PRIOR WORK

The development of geometric and surface shading descriptions and procedures for representing faces is known as facial modeling [Parke 1991]. The first computer-based three-dimensional character facial animation was generated by Parke at the University of Utah in the early seventies. By 1974, he had completed the first three-dimensional facial model that was parameterized (described in terms of a set of parameters) [Parke 1974]. Control parameter interpolation allowed manipulation of the model by an animator. Early parameter sets were primitive and developed with little attention to detailed facial anatomy, but they allowed a wide range of expressions for a wide range of facial conformations [Parke 1991]. A physically-based muscle-controlled facial model was introduced by Platt in the early eighties. It described facial expression through basic muscle actions modeled as springy meshes [Platt and Badler 1981]. In 1987, Waters combined ideas from these prior approaches with a new muscle-controlled facial expression model. This model used parameters to control underlying modeled musculature to enable more complex and realistic facial expression [Waters 1987].

So how do we capture the likeness of a known individual or even understand correct facial feature placement in virtual models? There are several sources of facial surface data; interactive sculpturing, assembly from components, creating new faces from existing faces, and three-dimensional surface measuring techniques. Measuring techniques include digitizing real surfaces by measuring the location of marked surface points on the surface to be measured, laser scanning of sculptures or real faces, and photogrammetric techniques (photographically capturing poses from several viewpoints) [Parke 1991]. Assembling faces from basic shapes like spheres, cones or cylinders can be useful for very simple characters. If we start with data for several polygonally modeled faces, we can use interpolation to create new faces by varying vertex positions between their various values from face to face [Parke 1972]. Finally, three-dimensional modeling programs with graphical user interfaces allow artists to create and manipulate surface geometry to create virtual face sculptures.

The work of Blanz and Vetter of the Max-Planck-Institute fur biologische Kybernetik in Tubinger, Germany is the primary motivation for my work. Following on early work by [Parke 1991], they began their work through the study of multi-dimensional interpolated facial spaces. Using linear combinations of a large number of face data sets, they defined multi-dimensional vector spaces to describe real faces. These data sets were obtained using a Cyberware laser scanner. These scans provided three dimensional facial surface data in a cylindrical coordinate system.

Their goal then was to derive a parametric description of all plausible faces by "exploiting the statistics of the data set." [Blanz and Vetter 2000] With this description, they were able to vary faces according to parameters associated with age, weight, sex, and features such as "hooked" noses. Their end goal was to represent "any face as a linear combination of a limited basis set of face prototypes." [Blanz and Vetter 2000] While their work focused on real faces, my work is focused on artistically created virtual face sculptures.

Principal component analysis, or PCA, of the data set is the primary tool that Blanz and Vetter used to define their facial parameters. Principal component analysis is a method of identifying patterns in data, and grouping this data in a way that identifies their similarities and differences [Hill and Lewicki 2006]. It is useful for analyzing data sets with high dimensionality that cannot be easily represented graphically. PCA begins by subtracting the mean value from each of the data dimensions. This produces a data set whose mean is zero. Next a covariance matrix is computed. Normally covariance is calculated between two dimensions, but a covariance matrix of all possible covariance values between all dimensions in the data set can be found as well. The eigenvectors and eigenvalues of this covariance matrix are used as the basis for analysis. PCA determines a linear transformation that transforms a data set to a new coordinate system such that the greatest variance of the data lies along the first coordinate axis (known as the first principal component), the second greatest variance lies on the second coordinate axis, and so on. PCA can be used to reduce the dimensionality of a data set by only keeping the most significant principal components and ignoring less important ones. Keeping only the most significant components means retaining those characteristics of the data set that contribute most to its variance. These components contain the most significant variation aspects of the data.

CHAPTER III

GOALS AND OBJECTIVES

There were several goals and objectives for this work. They are listed in the order in which in which they were addressed. These goals and objectives were to:

1. Develop a multi-dimensional vector space for describing faces starting from a set of exemplar three dimensional virtual face sculptures.
2. Identify and select the principal components of the exemplar data.
3. Reduce the dimensionality of the data by only retaining the most significant components. The significance of the components are determined by their corresponding eigenvalues.
4. Create a user interface for varying these components and graphically presenting the new faces created. This interface allows interactive manipulation of the principal component coefficients. The interface also enables the creation of animations showing the creation of new faces over time.
5. Explore how varying these principal components affect expression and facial features.
6. Create new faces based solely on the reduced dimension face space. This is done by varying only the retained components.
7. Extend the user interface to allow manipulation of characteristics such as gender, expression, and facial features. The basis for this control is difference vectors rather than principal components.

CHAPTER IV

METHODOLOGY

The approach is to develop and implement a prototype facial model based on analyzing data from artist created virtual facial sculptures. The facial sculptures are three dimensional surfaces created as versions of a conformable polygon mesh. This mesh was modified by 16 artists using standard interactive computer based three-dimensional surface sculpting tools. The artists were volunteers from the Texas A&M Visualization Sciences master's program. The initial provided mesh was a low polygon count surface (less than 1500 vertices) with head and facial features already generically modeled (Fig. 1). The artists were able to vary the positions of the vertices on this mesh, but were not be able to delete or add any edges, vertices, or polygons to the mesh. The mesh surface remained in its original three dimensional position and orientation. These restrictions allowed the face data to be represented as matrices with one to one vertex correspondence. Since the vertex topology and general facial features remained the same as the original mesh, all data gathered from various artists could be compared.

Each artist was asked to create a unique facial sculpture from this mesh in a neutral expression pose. The artists were then asked to create several versions of their face in different expression poses. They chose which expressions they would represent from the six universal emotions; fear, anger, surprise, joy, sorrow, and disgust. Each artist’s work was captured and exported as appropriate data files. The data was exported in the OBJ file format that could then be read into selected existing analysis software with minimal manipulation. PCA was then used to identify the principal components of this data.

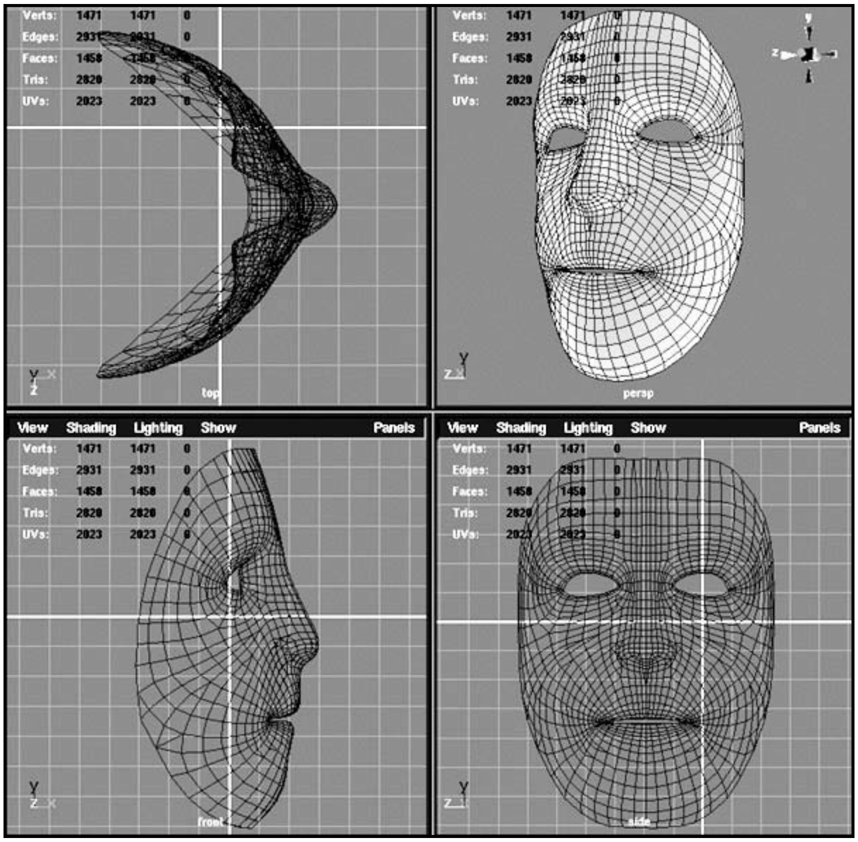

Figure1: The generic face mesh used.

Part of methodology included selecting analysis software best suited to the purpose and desired output for this research. Examples of known statistical packages that were applicable included Matlab, Scilab, and Praat which all include some form of PCA. I chose Matlab for my work. The statistical software was used to perform principal component analysis on the collected face data sets. A testing phase using a much simpler data set than the facial surface data was done to test the software and verify the technique before proceeding with the exemplar faces. This test data set was comprised of 13 mushroom shapes that I created from a generic mushroom mesh. The mushroom shapes contained fewer than ten percent of the number of vertices as the face sculptures. This made verifying the analysis process faster and simpler.

This research included developing interactive control of the derived principal components and graphically presenting their affects on the average faces. This was done by augmenting a pre-existing 3D modeling program with scripts, a series of commands, written in the program's native language. These scripts used the results calculated in the statistical software to modify and display the created faces. This user interface enabled the interactive creation of new face sculptures as well as animations showing linear interpolations between these created faces results over time.

A method for further modifying these newly created faces based on characteristics such as gender, facial feature and expression was explored through a second graphical user interface. This was done by modifying a created face along difference vectors associated with a specific characteristics. A difference vector is the vertex by vertex difference between two face data sets. In this work, a difference vector is between a weighted average of faces identified as having a specific characteristic, and the average face of the entire data set. Changing a face along a difference vector controlled the amount of influence the weighted characteristic had on the face.

CHAPTER V

IMPLEMENTATION

Exemplar Data Collection

First a polygonal mesh was created that formed a generic face. The goal was to create a generic facial mesh that would not influence the artist, such as by establishing gender or age or body weight. Another consideration was the density of the mesh. Although a larger number of vertices would allow more artistic freedom as well as the ability to create more detailed faces, a high density mesh would increase the complexity of the sculpting process and require greater computation time during statistical analysis and during use of a graphical user interface. The mesh used began as a mesh created by fellow graduate student Andy Smith based on a 3D scan of fellow graduate student Dave Morris' face (Fig. 2).

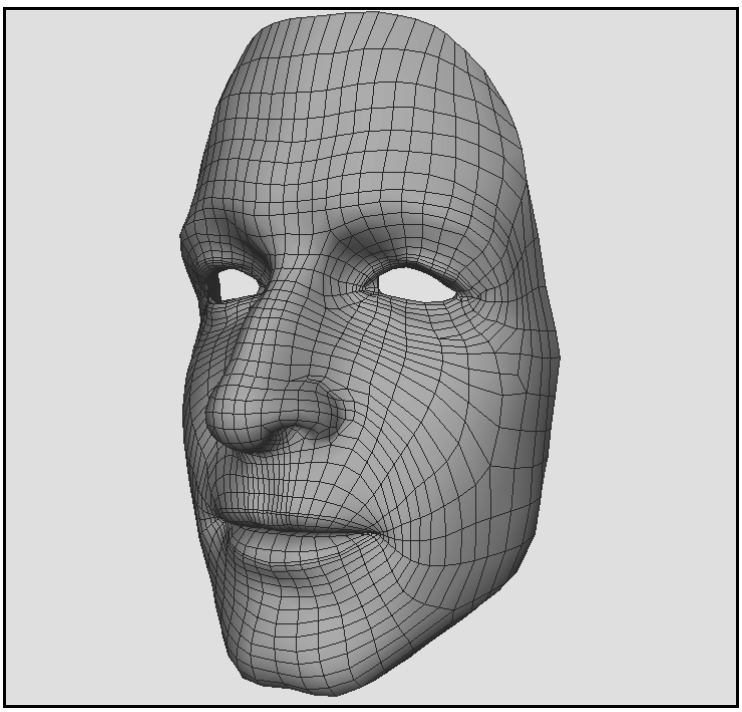

Figure 2: Mesh based on 3D scan of real face.

The topology of this mesh was altered by modifying unnecessarily dense areas. The mesh was then reshaped into a more generic face. Several versions of this generic mesh were created before arriving at the one finally used (Fig. 3). These versions vary slightly in their vertex count and triangulation of detailed areas like the corners of the mouth, eyes, and edge of the nose. The goal was to provide enough detail for the artists to work with while keeping the density of the mesh below 1500 vertices.

A digital copy of this mesh was distributed to each of the volunteer artists along with a copyright waiver in accordance to my Institutional Review Board approved protocol. After two months, a set of faces from each artists had been received. Each set consisted of the modified mesh in a neutral facial expression (Fig. 4) as well as three other expressions of the artist's choosing from the six universal emotions; fear, anger, surprise, joy, sorrow, and disgust. Figures 5 - 10 show the faces grouped by expression as identified by the artists. In these figures, the faces are all at the same scale. All of these sets together comprise the full data set of created exemplar faces (Fig.11).

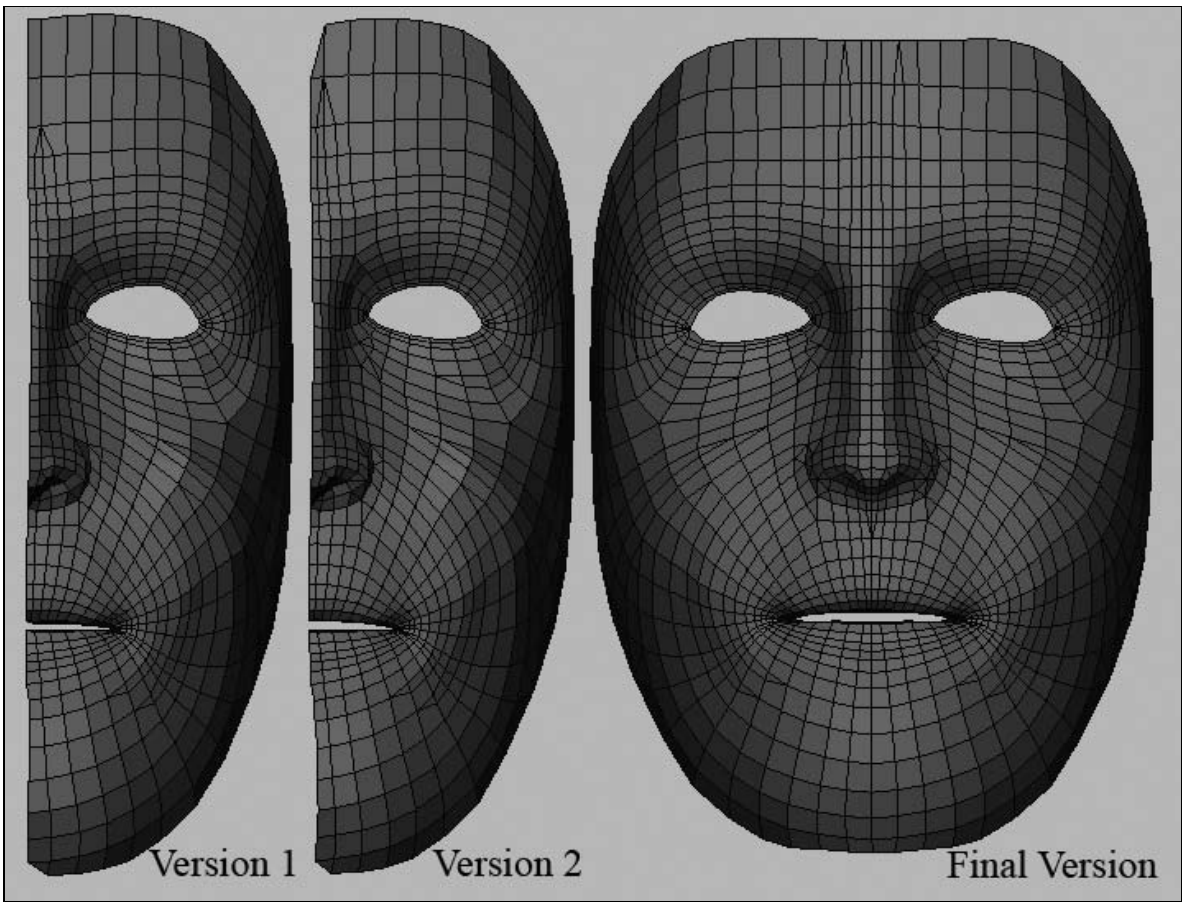

Figure 3: Stages of creating the original polygonal face mesh.

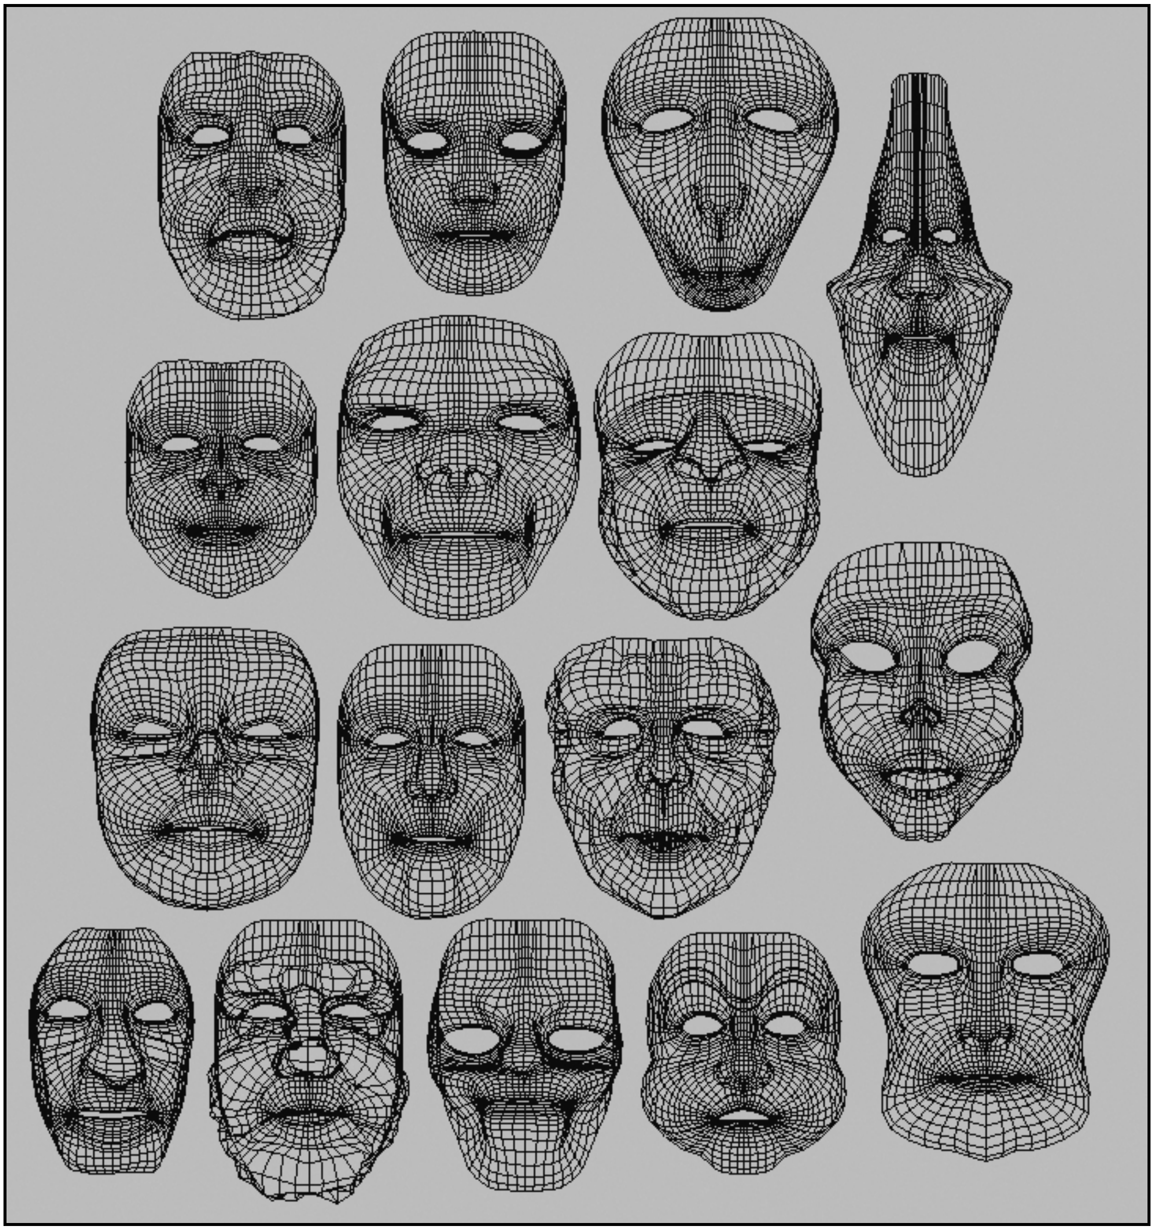

Figure 4: Neutral expressions in the face set.

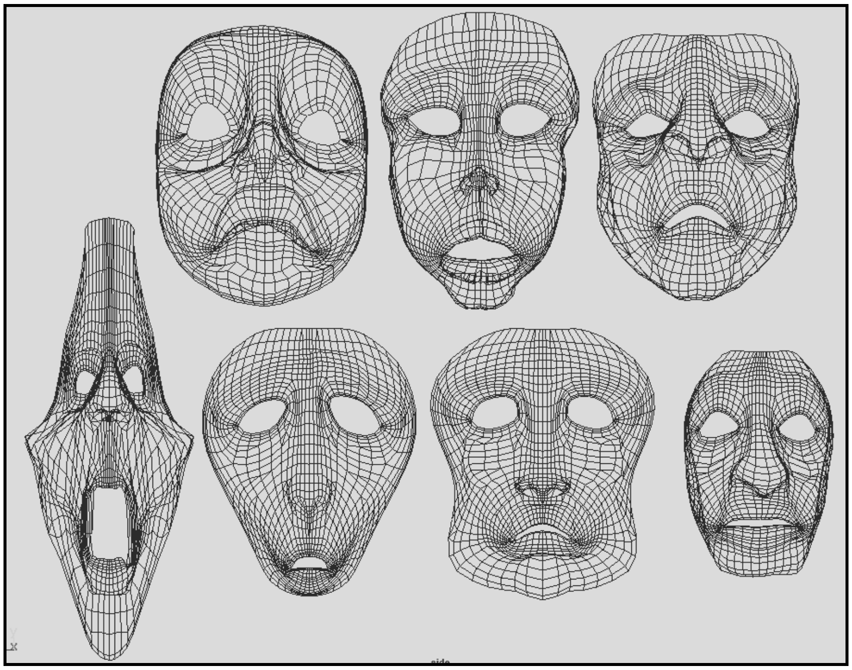

Figure 5: Fear expressions in the face set.

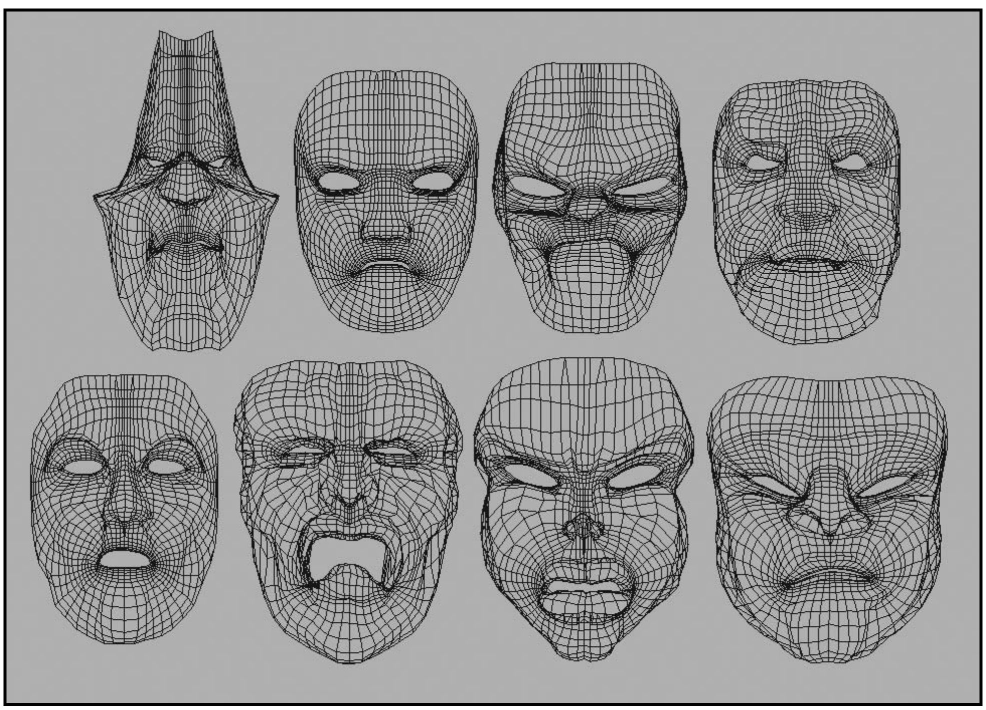

Figure 6: Anger expressions in the face set.

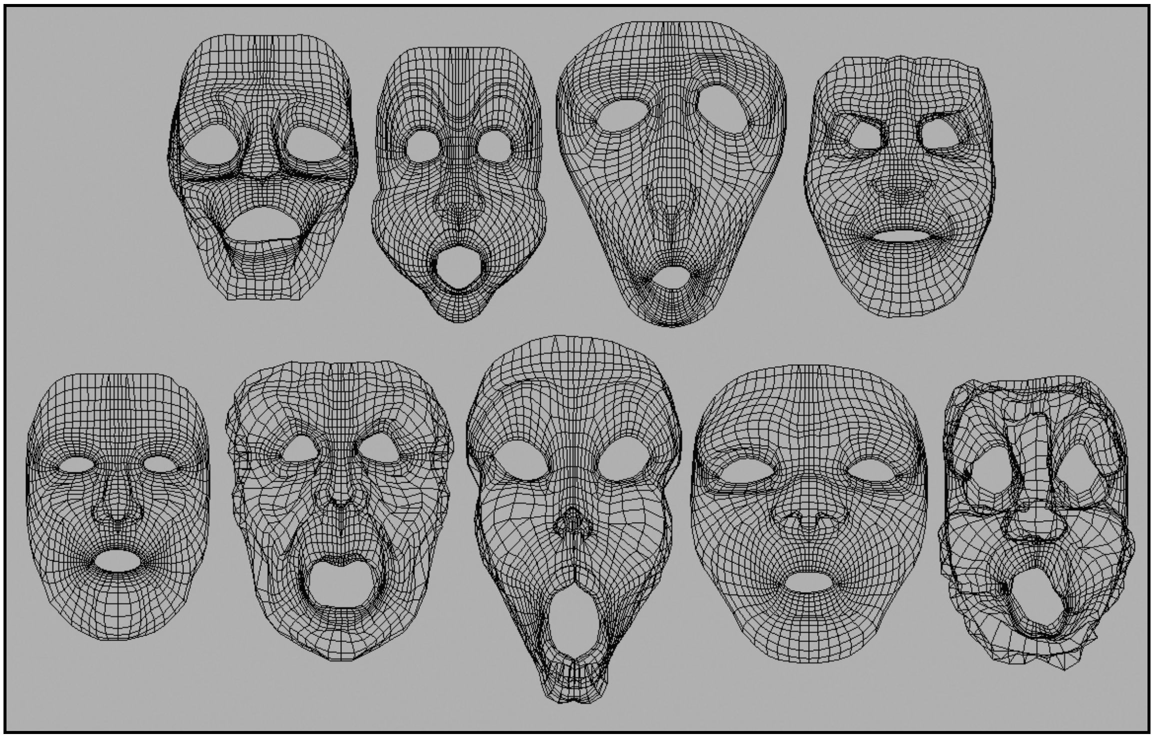

Figure 7: Surprise expressions in the face set.

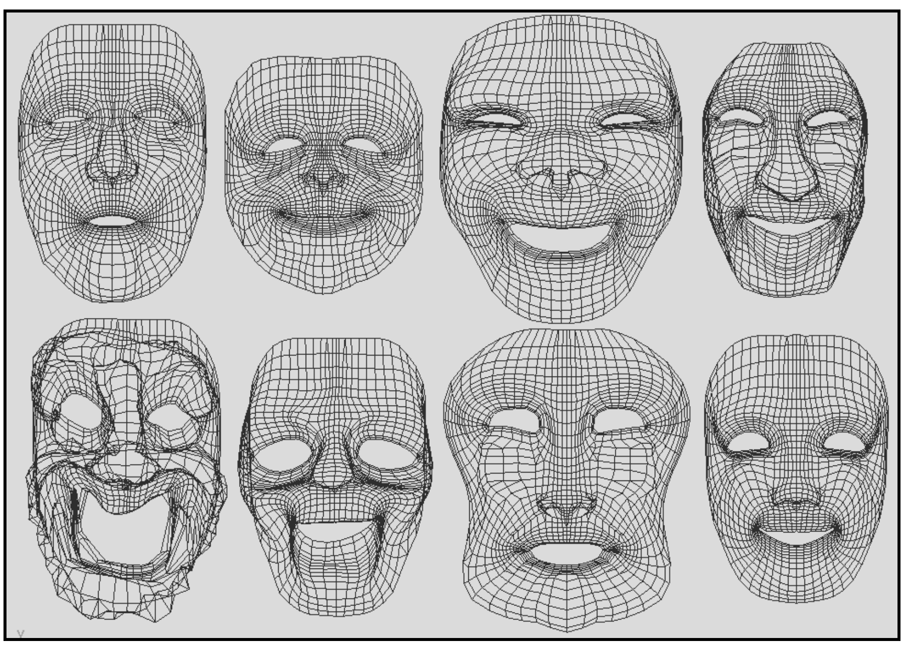

Figure 8: Joy expressions in the face set.

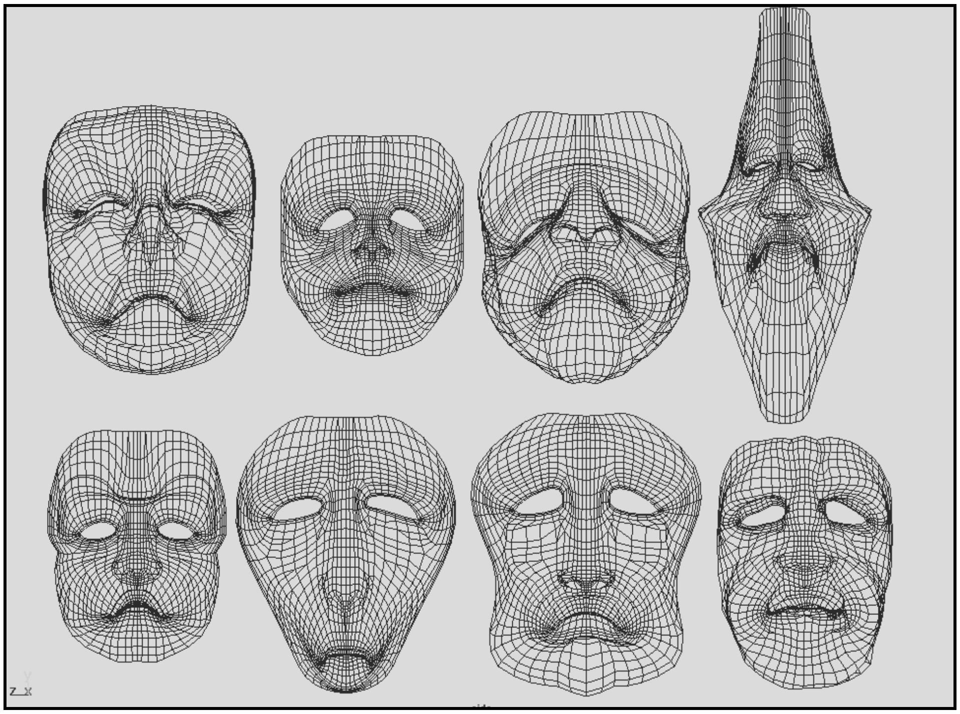

Figure 9: Sorrow expressions in the face set.

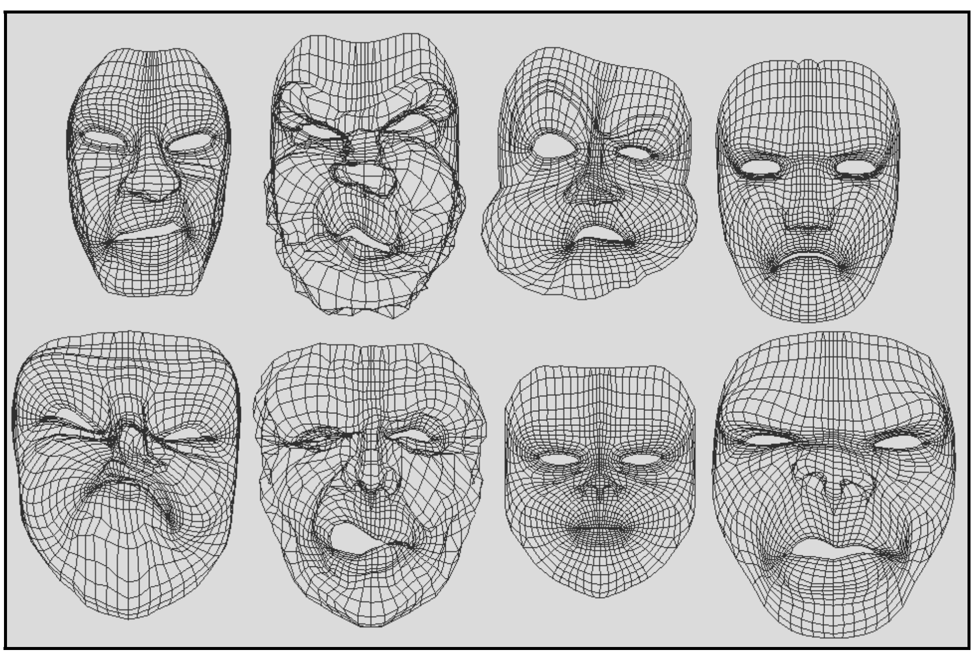

Figure 10: Disgust expressions in the face set.

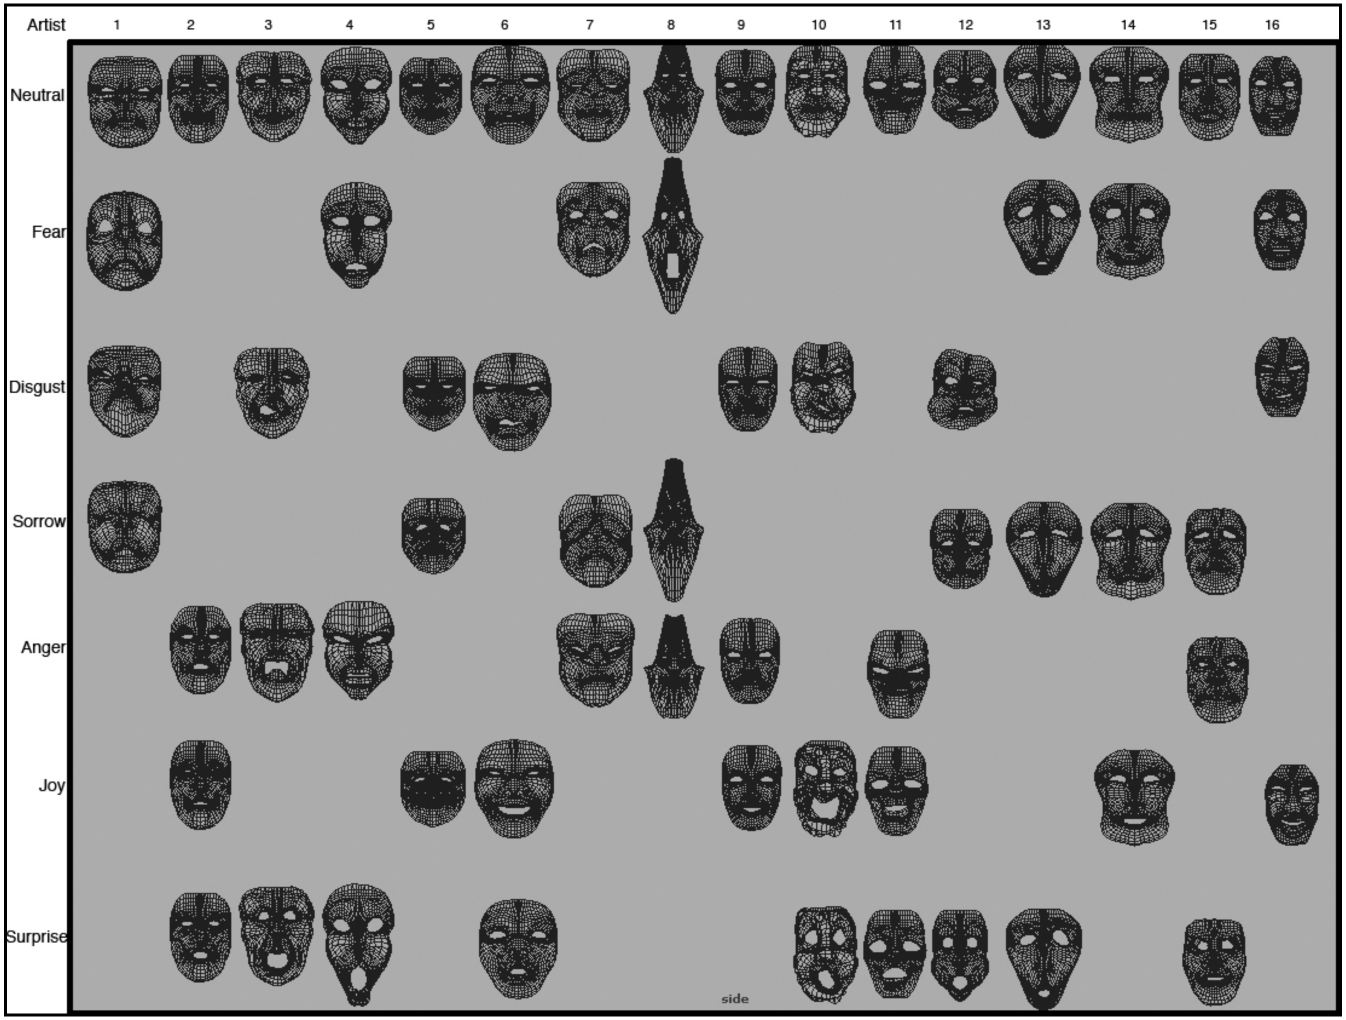

Figure 11: Entire face data set.

Testing Approach

Before working with the collected face data set, a testing phase with a much smaller data set was completed. The goal was to work through the analysis process with a test data set so that the process could be verified and problems could be resolved while working with fewer samples and a smaller number of vertices. For this a test mesh was needed which had enough vertices to define an appropriate 3D form, but which was simple enough to perform calculations on by hand if necessary for debugging and verification purposes. A basic mushroom shape was selected because it could be formed with fewer than 100 vertices and had a definite vertical orientation (Fig. 12). I then created 13 additional mushroom variations. These shapes varied in size, symmetry, and proportion. The test data set was formed of these 13 mushrooms shapes (Fig. 13).

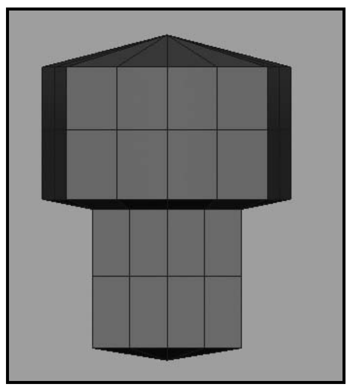

Figure 12: Original polygonal test mesh.

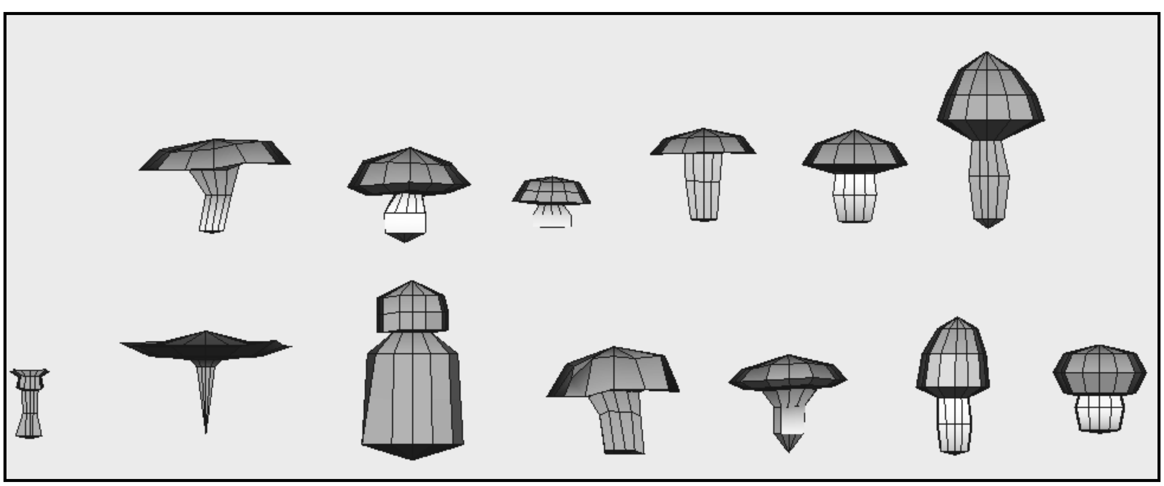

Figure 13: Entire test data set.

Test Data Preparation

Each mushroom mesh was exported as a text-based OBJ file from the 3D modeling software (Alias Maya) where they had been created. The OBJ format contains the vertex positions for an object in a sequence of x, y, and z coordinates, as well as corresponding texture and polygon vertex ordering information. An OBJ file saves the vertex coordinate information in vertex number order. The data is exported in the same order for each mushroom variation. This was key to my approach, as one-to-one vertex correspondence was needed between the different objects in the data set. Once exported each OBJ file had to be stripped of all non-numerical information such as alphabetical characters used as labels, as well as all the texture and polygon vertex ordering data. This stripped data was retained for later use. These stripped files were then used as input to the statistical software.

The position values of each vertex of every object in the data set needed to be stored such that the eigenvalues and eigenvectors of the covariance matrix could be calculated. Using a matrix, the vertex position values for the objects could be stored in columns, ordered by vertex number. In each column the x coordinates are followed by the y coordinates, and then the z coordinates. Each column represents an individual test object. In the case of the mushroom test data set, the matrix had 294 rows and 13 columns. (96 vertex positions, three rows each with an x-axis, y-axis, and z-axis coordinate, for each of 13 mushroom shapes.) Fig. 14 shows a small portion of this matrix. Even a mesh with fewer than 100 vertices requires a substantial matrix. The actual face data set was contained in a matrix with 4413 rows and 64 columns.

Data Analysis

The Mathworks' software Matlab was chosen as the statistical tool to be used because of its availability and built-in principal component analysis functions. The data in the mushroom shape files were used to form a matrix containing the entire test data set vertices. There were two main analysis objectives; to calculate an average mushroom mesh from the 13 created mushroom meshes, and to perform principal component analysis. The first objective was met by finding the average coordinates for each vertex of the mesh. The second objective required principal component analysis on the test data set matrix. The statistical software calculated the eigenvectors (as a matrix) and the corresponding eigenvalues of the data (Fig. 15). The eigenvector whose eigenvalue is the largest is the first principal component. Other principal components are ranked based on their eigenvalues.

```
>> loadfilesMush(names)

Vertices =

   -0.4181   -0.4293   -0.3255   -0.5000   -0.5000   -0.3041   -0.4908   -0.3435   -0.3471   -0.4089   -1.0719   -0.2319   -0.5000
         0         0   -0.0006         0         0   -0.0007         0         0    0.0002   -0.0007         0         0         0
    0.4181    0.4293    0.3244    0.5000    0.5000    0.3028    0.4908    0.3435    0.3475    0.4076    1.0719    0.2319    0.5000
   -0.3763   -0.5000   -0.4930   -0.3142   -0.2939   -0.5995   -0.5332   -0.3462   -0.5938   -0.7359   -0.4495   -0.8918   -0.2156
    0.0006         0   -0.0000    0.0006    0.0011    0.0004    0.0001   -0.0007    0.0003    0.0010   -0.0002    0.0016   -0.0012
    0.3776    0.5000    0.4929    0.3155    0.3239    0.6003    0.5334    0.3449    0.5944    0.7379    0.4491    0.8951    0.2133
   -0.4181   -0.4293   -0.3255   -0.5000   -0.5000   -0.3041   -0.4908   -0.3435   -0.4412   -0.4089   -1.3430   -0.2319   -0.5000
         0         0   -0.0006         0         0   -0.0007         0         0    0.0006   -0.0007         0         0         0
    0.4181    0.4293    0.3244    0.5000    0.5000    0.3028    0.4908    0.3435    0.4423    0.4076    1.2676    0.2319    0.5000
    0.3776    0.5000    0.4929    0.3155    0.3239    0.6003    0.5334    0.3449    0.7204    0.7379    0.4491    0.8951    0.2133
   -0.3763   -0.5000   -0.4930   -0.3142   -0.2939   -0.5995   -0.5332   -0.3462   -0.7186   -0.7359   -0.4495   -0.8918   -0.2156
   -1.2366   -0.5823   -0.5741   -0.8381   -1.1387   -0.6983   -1.1169   -0.6928   -0.7963   -0.8573   -0.6026   -1.1963   -0.2887
    0.0019   -0.0003   -0.0003    0.0006    0.0205    0.0000    0.0001    0.0001    0.0003    0.0005   -0.0002    0.0016   -0.0012
    0.0012   -0.1176    0.0415    0.0001    0.2667    0.0025    0.0008   -0.0002    0.0008    0.0035   -0.0002    0.0036   -0.0013
   -1.0516   -2.3187   -1.1240   -0.7062   -1.5522   -1.4642   -1.3270   -0.6001   -0.9572   -1.7981   -0.6026   -1.8311   -0.2395
    1.2404    0.5817    0.5735    0.8394    0.9306    0.6983    0.7151    0.7506    0.7969    0.8583    0.7539    1.1995    0.2864
    1.0539    1.1268    1.1268    0.7064    1.5957    1.3717    1.2566    0.6572    0.9588    1.6855    0.7539    1.8383    0.2368
   -1.3074   -1.3974   -2.2074   -0.8778   -1.4976   -1.7002   -1.1499   -0.7458   -1.4404   -2.0880   -0.8919   -2.9877   -0.2973
    1.3032    1.3933    1.3933    0.8737    1.4740    1.6960    1.1458    0.7993    1.4363    5.6064    0.8966    2.9445    0.2932
    1.5327    0.7191    0.7089    1.0374    1.0642    0.8632    0.8839    0.9142    1.1933    1.0608    0.8961    1.4823    0.3543
   -1.5375   -0.7237   -1.5235   -1.0419   -1.1001   -0.8679   -0.8886   -0.8611   -1.1981   -1.0657   -0.7490   -1.4873   -0.3586
   -0.6456   -1.8868   -0.8352   -0.5799   -1.1141   -1.0879   -1.1180   -0.3992   -0.7963   -1.3359   -0.6026   -1.1963   -0.2810
   -0.0001   -0.1189    0.0406   -0.0003    0.0278    0.0013    0.0001   -0.0008    0.0003    0.0021   -0.0002    0.0016   -0.0012
    0.0006   -0.1175    0.0014   -0.0003    0.0067    0.0022    0.0001   -0.0008    0.0009    0.0032   -0.0002    0.0016   -0.0012
   -0.8024   -2.1905   -1.8483   -0.7208   -1.0831   -1.2631   -0.8900   -0.4960   -1.1981   -1.5512   -0.7490   -1.4873   -0.3490
    0.6454    2.0876    0.8363    0.5794    1.1724    1.0182    1.0463    0.3975    0.7969    1.2512    0.6022    1.1995    0.2786
    0.7983    2.8014    1.0342    0.7166    1.0782    1.2590    0.8859    0.4918    1.1940    5.0695    0.7449    1.4832    0.3449
   -0.5546   -0.5695   -0.4316   -0.5000   -0.5000   -0.2969   -0.6064   -0.4278   -0.3471   -0.3992   -0.9622   -0.4128   -0.1798
         0         0   -0.0006         0         0   -0.0007         0         0    0.0002   -0.0007   -0.0040    0.0016         0
    0.5546    0.5695    0.4305    0.5000    0.5000    0.2956    0.6064    0.4278    0.3475    0.3978    0.9542    0.4160    0.1798
    0.5546    0.5695    0.4305    0.5000    0.5000    0.2956    0.6064    0.4278    0.4423    0.3978    1.1291    0.4160    0.1798
   -0.5546   -0.5695   -0.4316   -0.5000   -0.5000   -0.2969   -0.6064   -0.4278   -0.4412   -0.3992   -1.2045   -0.4128   -0.1798
   -0.3228   -0.7440   -0.4174   -0.2901   -0.5011   -0.5433   -0.7599   -0.2000   -0.3980   -0.6669   -0.3014   -0.5973   -0.1411
   -0.5252   -1.4770   -0.5613   -0.3530   -0.4773   -0.7309   -0.8641   -0.3002   -0.4782   -0.8973   -0.3014   -0.9138   -0.1204
   -0.6174   -0.2913   -0.2872   -0.4187   -0.5302   -0.3491   -0.7594   -0.3463   -0.3980   -0.4284   -0.3014   -0.5973   -0.1449
   -0.1878   -0.2500   -0.2465   -0.1568   -0.1603   -0.2995   -0.2665   -0.1734   -0.2968   -0.3675   -0.2249   -0.4451   -0.1084
   -0.2773   -0.2847   -0.2161   -0.2500   -0.2500   -0.1488   -0.7498   -0.2139   -0.1735   -0.1999   -0.4831   -0.2056   -0.0899
   -0.2091   -0.2147   -0.1630   -0.2500   -0.2500   -0.1524   -0.2454   -0.1717   -0.1735   -0.2048   -0.5360   -0.1159   -0.2500
    0.3227    0.5062    0.4184    0.2895    0.6001    0.5096    0.3582    0.1983    0.3986    0.6265    0.3010    0.6006    0.1387
    0.5276    0.5641    0.5641    0.3533    1.0495    0.7358    0.4637    0.2997    0.4798    0.9043    0.3010    0.9086    0.1177
    0.6211    0.2907    0.2866    0.4200    0.5886    0.3492    0.3576    0.3466    0.3986    0.4294    0.3010    0.6006    0.1426
    0.1891    0.2500    0.2464    0.1581    0.1625    0.3003    0.2668    0.1721    0.2974    0.3694    0.2244    0.4483    0.1061
    0.2773    0.2847    0.2150    0.2500    0.2500    0.1474    0.3032    0.2139    0.1738    0.1986    0.4751    0.2088    0.0899
    0.2091    0.2147    0.1619    0.2500    0.2500    0.1511    0.2454    0.1717    0.1738    0.2035    0.5360    0.1159    0.2500
   -0.4181   -0.4293   -0.3255   -0.5000   -0.5000   -0.3041   -0.4908   -0.3435   -0.4412   -0.4089   -1.3430   -0.2319   -0.5000
   -0.5546   -0.5695   -0.4316   -0.5000   -0.5000   -0.2969   -0.6064   -0.4278   -0.4412   -0.3992   -1.2045   -0.4128   -0.1798
   -0.3763   -0.5000   -0.4930   -0.3142   -0.2939   -0.5995   -0.5332   -0.3462   -0.7548   -0.7359   -0.4495   -0.8918   -0.2156
   -1.3871   -0.6530   -0.6438   -0.9400   -1.0660   -0.7831   -0.8018   -0.7770   -1.1354   -0.9615   -0.6758   -1.3418   -0.3237
   -1.1795   -1.6068   -1.2607   -0.7920   -1.7559   -1.5338   -1.0375   -0.6730   -1.3650   -1.8837   -0.8187   -2.2754   -0.2684
   -0.7240   -2.0387   -0.9368   -0.6504   -1.0468   -1.1396   -0.8030   -0.4476   -1.1354   -1.3994   -0.6758   -1.3418   -0.3150
    0.7218    2.2356    0.9352    0.6480    1.0542    1.1386    1.1311    0.4447    1.1342    4.9216    0.6735    1.3414    0.3117
    1.1786    1.2600    1.2600    0.7900    1.4408    1.5339    1.3662    0.7282    1.3645    5.4072    0.8253    2.4671    0.2650
    1.3866    0.6504    0.6412    0.9384    1.0402    0.7808    0.7995    0.8324    1.1338    0.9596    0.8250    1.3409    0.3204
```

Figure 14: Screen capture of a portion of the test data set matrix.

```
>> loadfilesMush(names)

Eigenvalues =

  199.2558
   13.6859
    5.4953
    3.4252
    2.8116
    1.6342
    1.3069
    0.5570
    0.4135
    0.2857
    0.1819
    0.0885
```

Figure 15: Screen capture of the eigenvalues for the first test data set.

For viewing, the average vertex positions were exported as a text file. This data was used to reconstruct a valid OBJ file. All the meshes have identical polygon ordering (texturing information was not used) so the polygon ordering data from any of the original created mushroom mesh OBJ files could be used. This new OBJ file could be viewed using existing 3D software such as Alias Maya. The average mushroom shape could then be compared to the original mushroom mesh and to the created mushrooms (Fig. 16). Because this average mushroom appeared perfectly symmetrical, the test process was repeated with some new created mushroom shapes that had more extreme forms; greater asymmetry, size difference, and inverted proportions such as bases larger than their tops. The average mushroom from this second data set reflected these changes (Fig. 17).

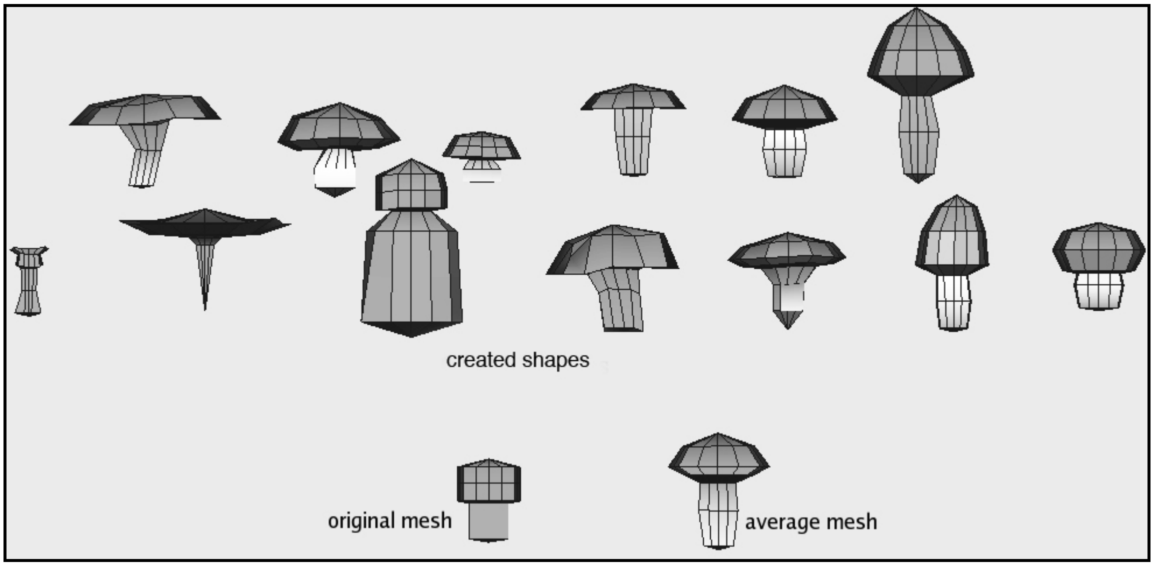

Figure 16: Comparing the original and average test meshes for the first test set.

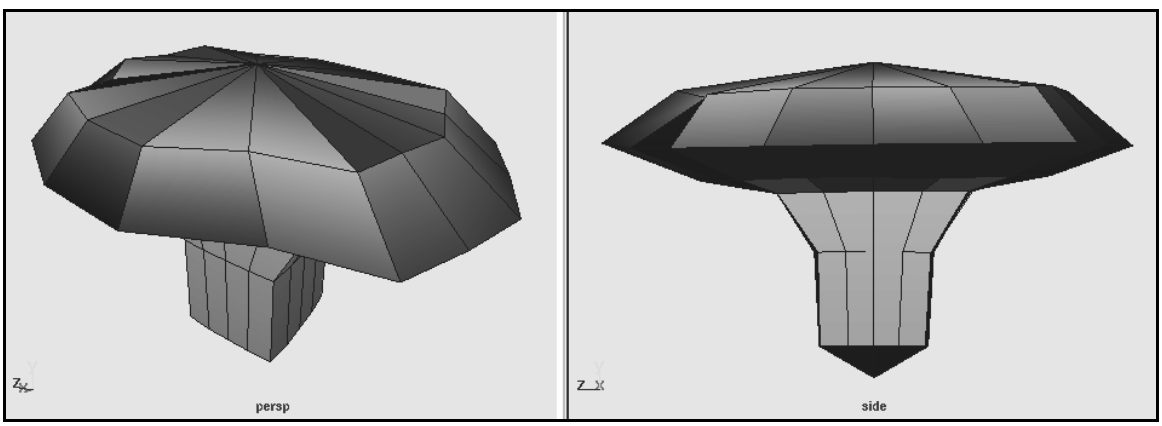

Figure 17: Average mesh from the second test data set.

Calculating the principal components of the test data set allowed creating new mushroom shapes in the PCA-based shape space. A multi-dimensional vector space for describing mushrooms was created. This space was such that the greatest variance by any projection of the mushroom data in the Cartesian coordinate system lay along the first coordinate of the new space. The second greatest variance of the data set lay on the second coordinate, and so forth. These coordinates are also known as the principal components, or eigenvectors, of the data set. The dimensionality of the new vector space can be reduced by keeping only the most significant principal components in the coordinate system and deleting less significant components. The retained principal components correspond to the most significant variation of the data set while the deleted components correspond to less significant variation. The significance of the components are determined by their eigenvalues. These eigenvalues correspond to the amount of variance of the data along the corresponding eigenvector. The largest eigenvalue corresponds to the eigenvector known as the first principal component. The number of non-zero eigenvalues is equal to one less than the number of linearly independent vectors in the original data matrix. In this case, each data set had one fewer non-zero eigenvalues than there were objects (faces or mushrooms) in the set.

This change of basis allowed me to create new mushroom shapes based on a reduced dimensionality space, retaining only the first and second principal components. These new mushroom shapes were created by varying the influence of only the first two components. For each principal component, its eigenvector, scaled by a coefficient, is added to the average mushroom to morph (to change form or shape of) the average mushroom along the corresponding principal component. The coefficient value determines the amount of influence that eigenvector, or component has on the new mushroom shape. Several example created mushroom shapes were tested and viewed in Maya. I then felt confidant that I could repeat this process with the collected face data set.

Face Data Analysis

All the face files had to be name coded for privacy in accordance with the Institutional Review Board approved protocol. They were then exported as OBJ files and formatted for use by the statistical software. The mushroom test data set only required the calculation of one average mesh. However, for the face data set I calculated the average face from all faces in the data set (Fig. 18), as well as the average face for each of the expression sub-sets (Fig. 19). The eigenvectors and eigenvalues for the entire set and for each subset were calculated in the same manner as with the test data set.

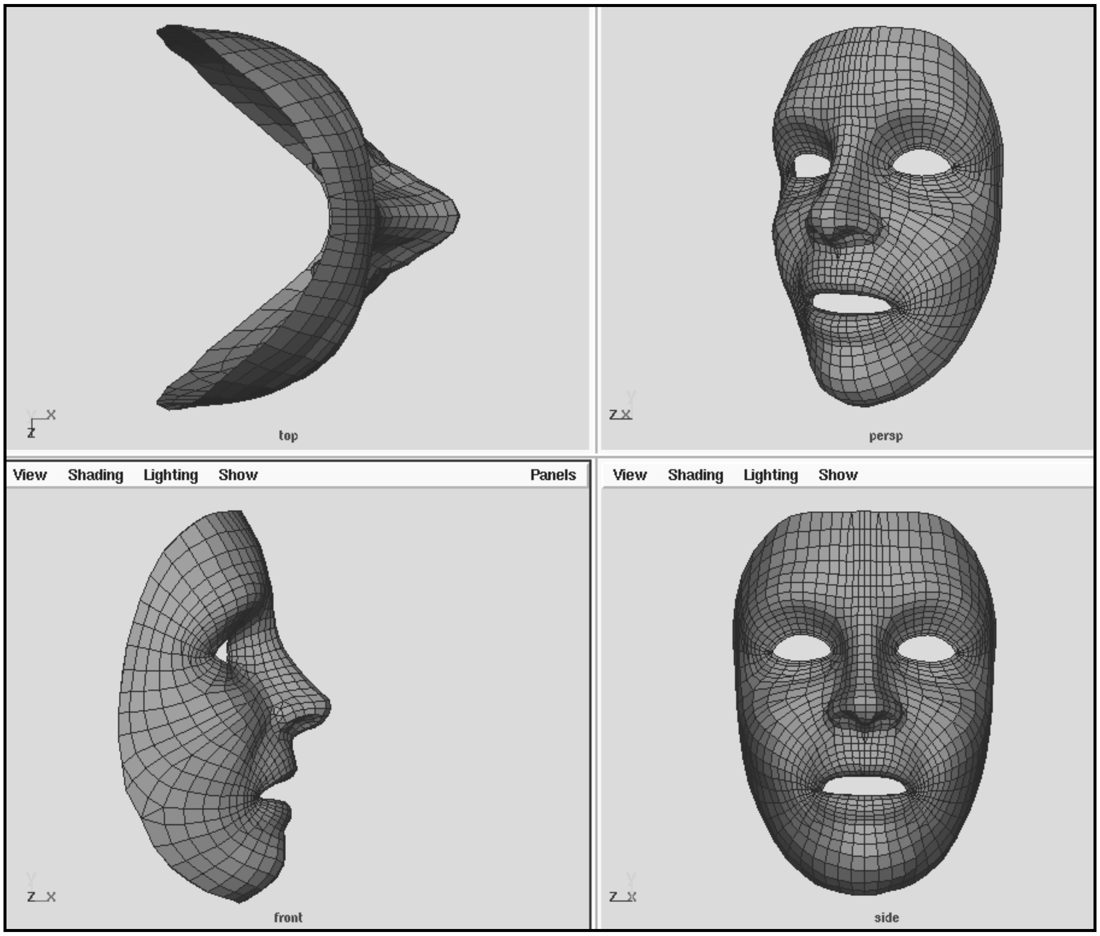

Figure 18: Views of the average facial mesh from the entire face data set.

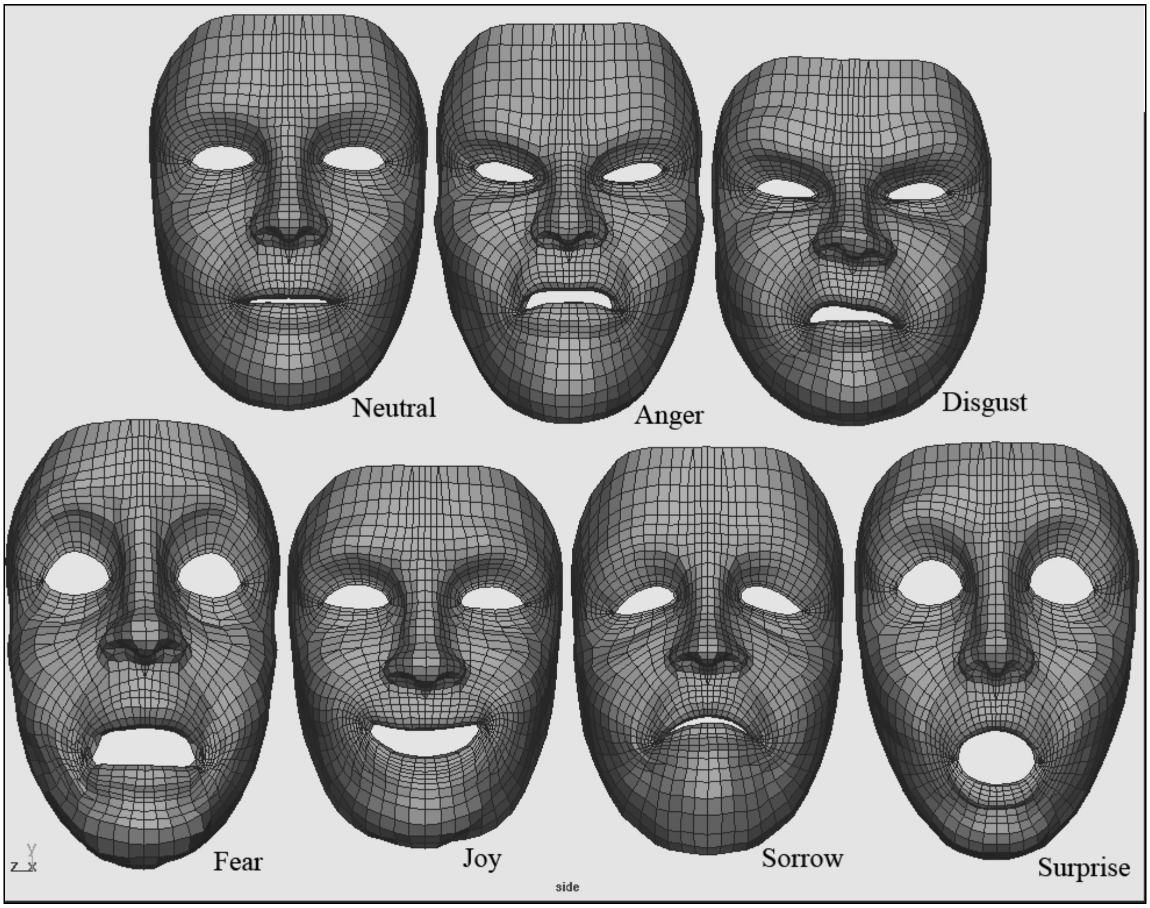

Figure 19: Average facial meshes from each expression data set.

Graphical User Interfaces

For the test data, the newly created mushroom shapes calculated from Matlab had been formatted by hand and brought into 3D modeling software for viewing. The next goal was to create a graphical user interface for creating new faces that would allow combining variably scaled eigenvectors, or principal components. This would enable interactive exploration of the principal component based face space. The easiest way to create such an interface was to take advantage of a pre-existing 3D modeling software's ability to display polygonal meshes and to manipulate those meshes interactively. An interface was desired that would allow the user to interactively change the coefficients of the eigenvectors being added to the average face mesh and to see the mesh updated quickly. The interface also needed to handle all of the desired principal components at once. Alias' Maya 3D modeling program was used to create the graphical interface. Maya provides a scripting language, MEL, useful for developing such interfaces. By adding interface scripts to Maya, it could be used to control and display the created facial meshes. These scripts enabled me to use the results of calculations performed in the statistical software to manipulate and display facial meshes.

Multiple Maya scene files, which contain the needed surface geometry, were created that each showed either the average facial mesh for the entire face data set or one of the average facial meshes for the expression sub-sets; joy, fear, anger, sorrow, surprise, and disgust. This allowed the creation of separate interfaces for each of the sets. These ‘average’ meshes held the vertex positions from the statistical calculations. Because the scaled eigenvectors would be added to the average vertex coordinates, the average positions needed to remain fixed. Once manipulated, the facial mesh would no longer hold the original vertex coordinates. Each subsequent change to the eigenvector coefficient or the addition of multiple scaled eigenvectors would further distort the vertex positions. To resolve this, in each Maya scene file, the interface script accessed average vertex position data from a hidden (non-displayed geometry data) mesh while manipulating a separate visible mesh. This retained the average mesh vertex positions regardless of how the visible mesh was manipulated.

The calculated eigenvalues and eigenvectors had been saved from the statistics program as text files which could then be accessed from the 3D modeling program. Maya’s Embedded Language (MEL) has commands for accessing a specific vertex position vector by vertex number, as well as for defining new values for those vectors. These commands were used in the interface scripts to add the scaled eigenvectors to the average facial mesh's vertex positions. Each value in an eigenvector corresponds with a specific vertex x, y, or z coordinate. In this way the interface could manipulate the mesh according to the scaled values for each eigenvector.

To allow the user to manipulate these values, a screen window of value sliders was created. These sliders could adjust values that were associated with the various principal component coefficients (Fig. 20). Each time a slider was changed, the face mesh would interactively update its shape (Fig. 21-22). These sliders could be used individually, or together to create new faces (Fig. 23). The slider values of zero in Figure 20 would result in the average face without any variation from the data set's principal components.

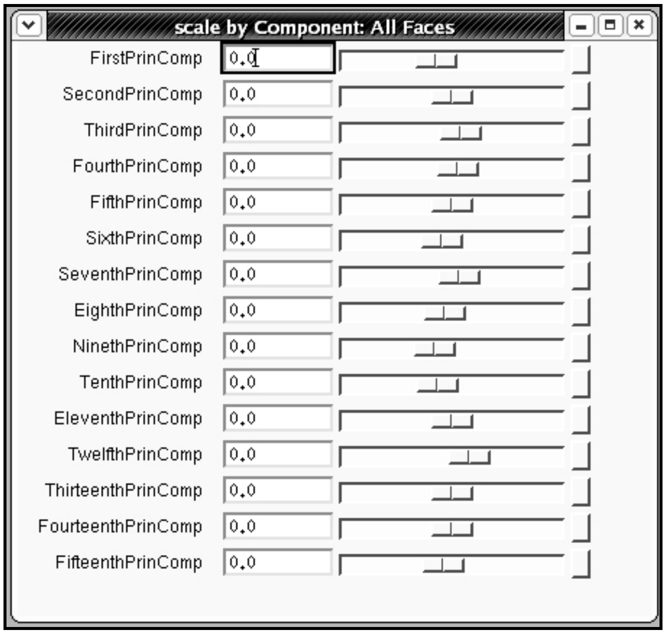

Figure 20: Graphical user interface sliders to control the value of component coefficient.

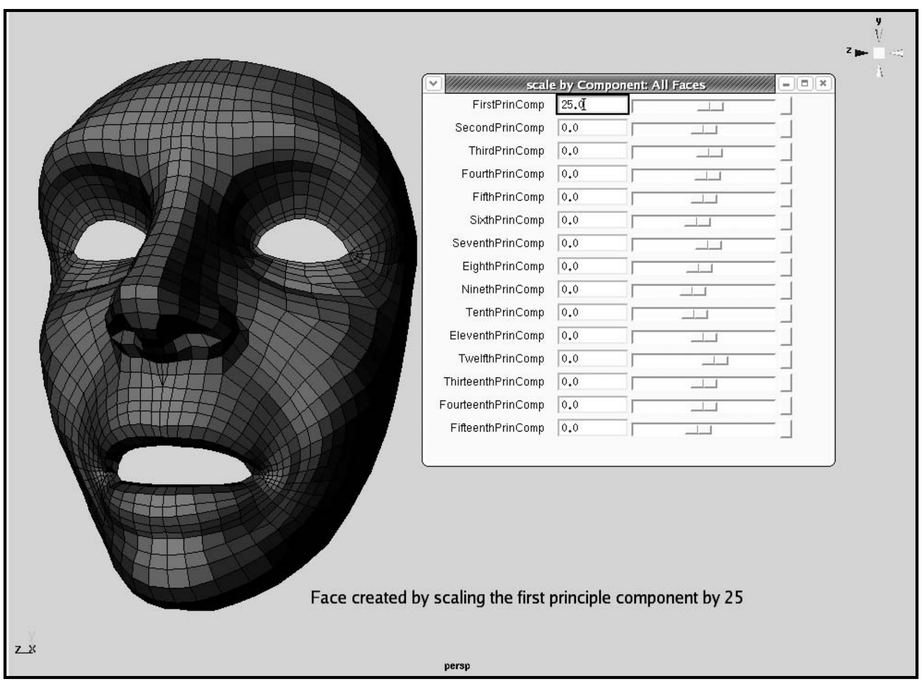

Figure 21: Modifying the first principal component coefficient using the user interface.

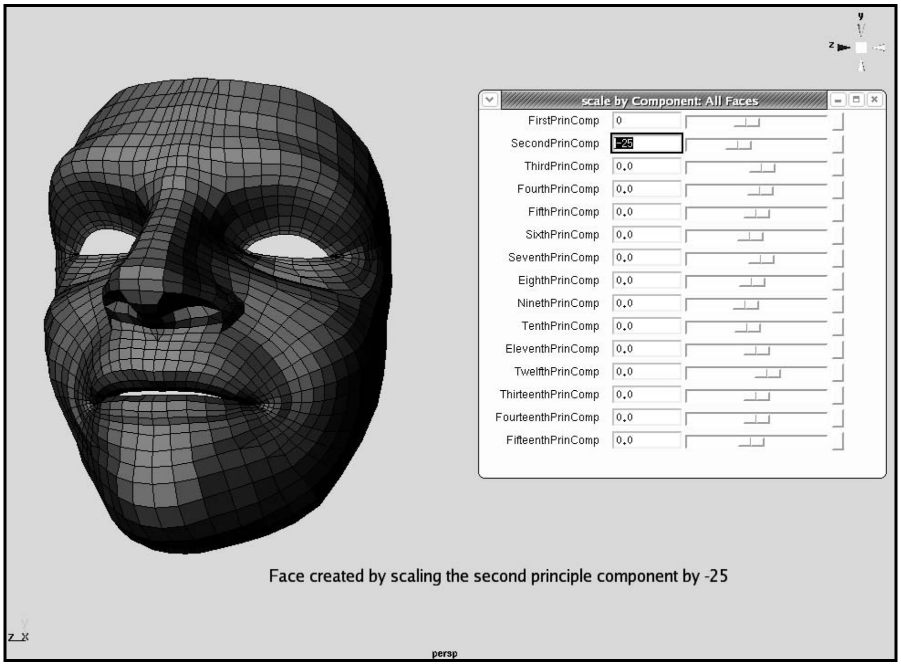

Figure 22: Modifying the second principal component coefficient using the user interface.

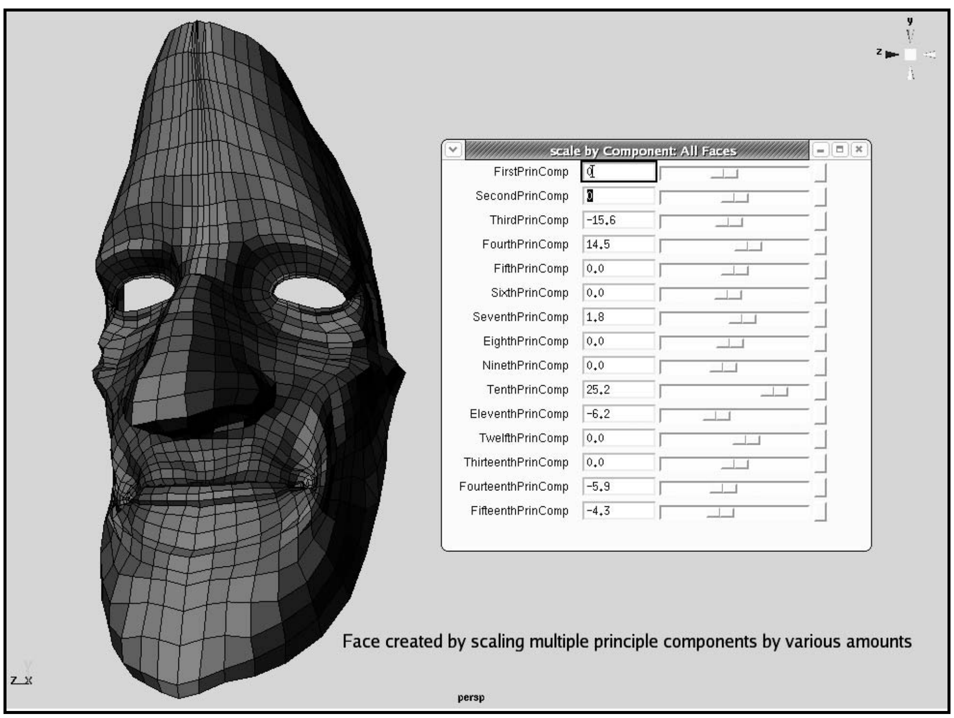

Figure 23: Modifying multiple component coefficients using the user interface.

Keyable Interface

The interface described above worked very well for exploring manipulation of the face mesh. But it did not allow the creation of animations showing changes in facial form over time. Animation required interpolation between key poses of the mesh. The initial interface could not save poses of the mesh because Maya's pre-existing animation control interface could not identify the variables needed to save the pose of the mesh. An alternative interface was needed. This was done by creating scene files such that the average mesh was associated with Maya attributes that corresponded to, and controlled the value of, each of the principal component coefficients. These attributes are properties of the mesh and can be seen in the right of Figure 24. A new user interface was created where the sliders controlled the values of these attributes instead of directly altering the component coefficients. The values of these attributes could be saved like any other keyable attribute, such as position or scale of an object. With this interface animations could be created of the mesh morphing between sets of saved coefficient values.

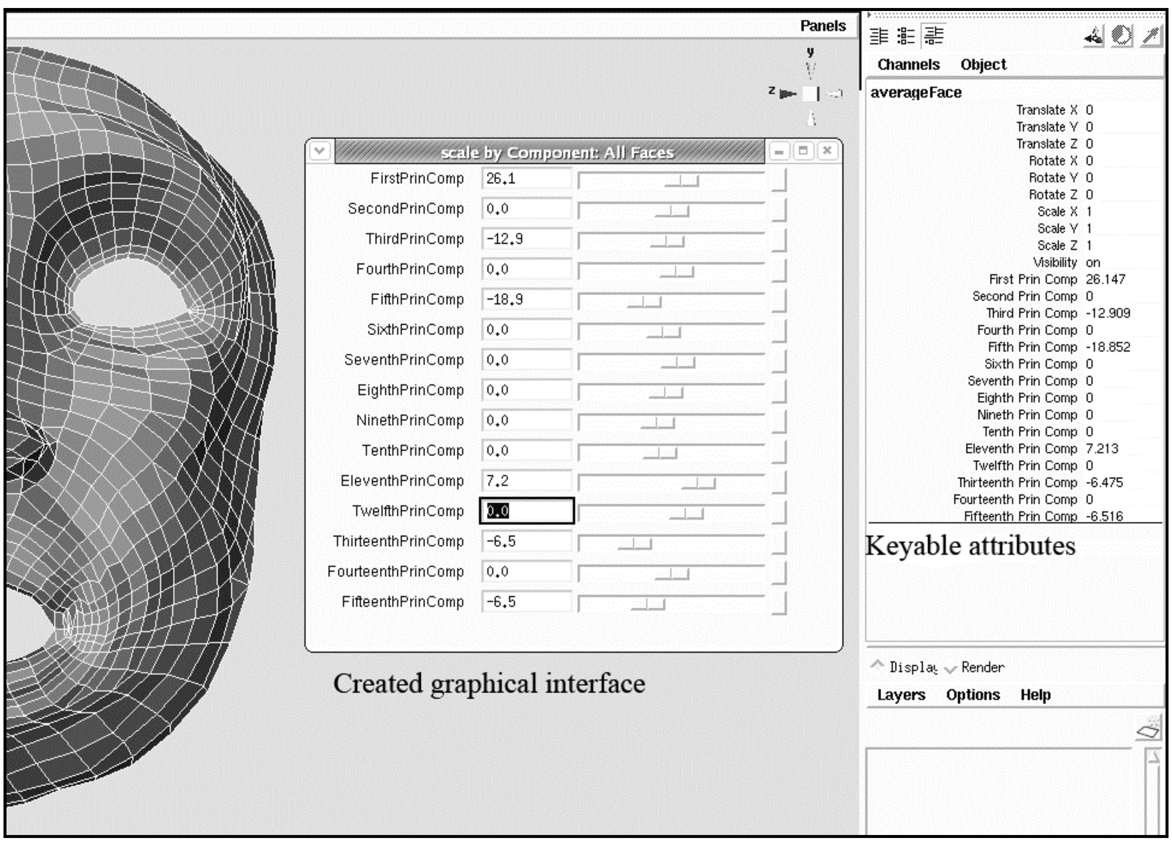

Figure 24: Controlling keyable attributes using the user interface.

Feature-based Faces

I also wanted to explore a technique and associated interface that would influence a created face's form not through its principal components, but through specific characteristics such as expression, gender, or facial features. Once a new face was created using the principal component interface it could then be further manipulated according to these characteristics. This was done using difference vectors. For each chosen characteristic, the average face mesh of the data set and a weighted average of all faces in the data set containing the specific characteristic are used. The weighted average face is such that it contains the maximum variance of the selected characteristic. A difference vector is a vector representing the difference between the average face mesh and a weighted face mesh. Changing a face along the difference vector controls the amount of influence that characteristic has on the created face. Using the same approach used with the principal components, I created a slider interface for varying the coefficient of each difference vector (Fig. 25).

A weighted average was calculated of faces in the entire face data set based on the appearance of a specific feature. For instance, for the joy expression, the most joyous face would be weighted as 1.0, while a face of lesser joy could have a weighting of 0.2. A non-joyful expression would be weighed as zero, effectively leaving it out of the calculation. Subtracting this weighted average from the average joy face produced a difference vector. This vector could be used to vary the amount of joyful expression of a facial mesh (Fig. 26). Gender as a variable attribute was also explored (Fig. 27). Specific facial features could be used as well. To manipulate the nose of the mesh, facial meshes in the data set whose nose turned upwards were weighed as positive. Facial meshes whose noses were down-turned were weighed as negative. The end result was a control that could turn the created face's nose more up or down (Fig. 28).

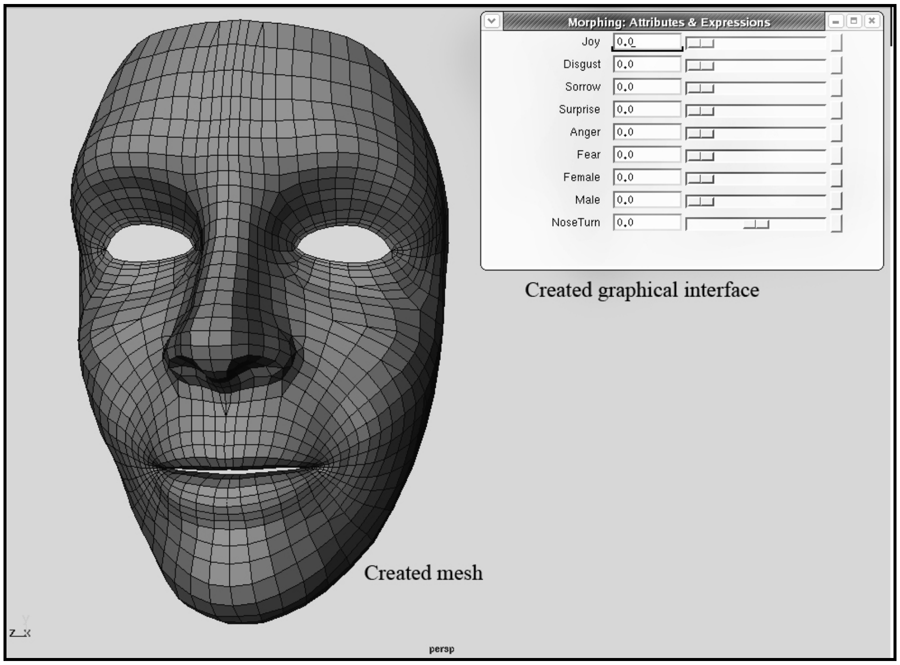

Figure 25: User interface extended to control specific attributes.

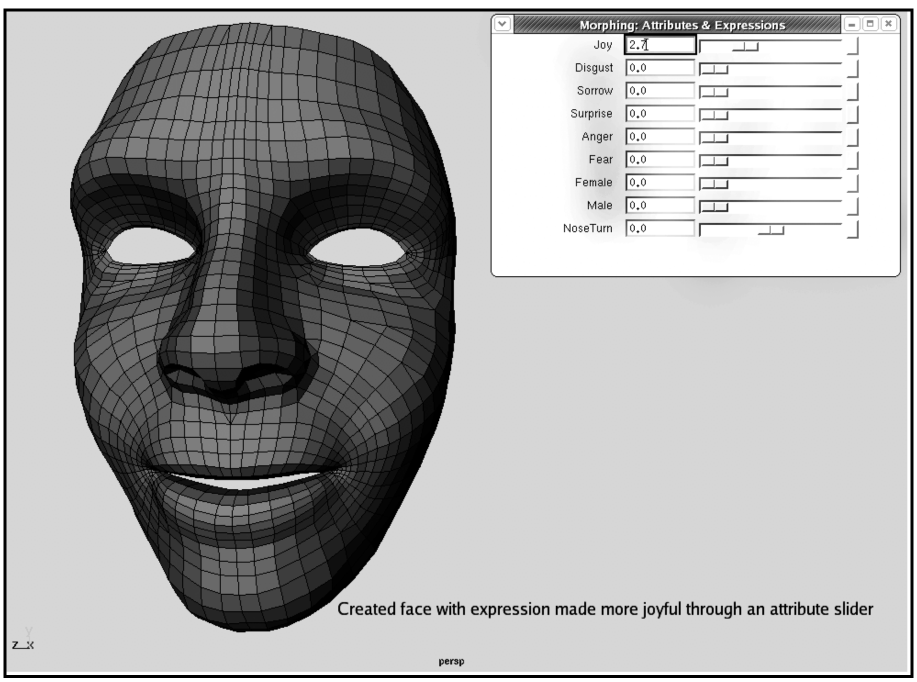

Figure 26: Increasing the joy expression of a created facial mesh.

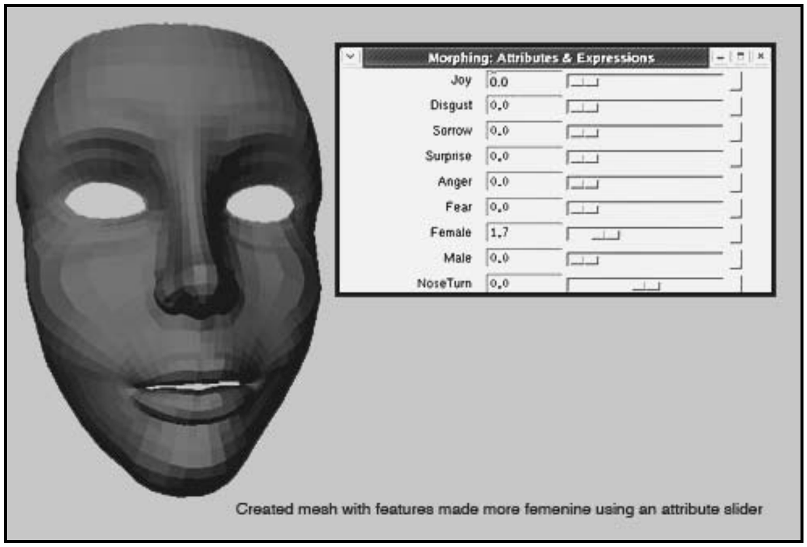

Figure 27: Increasing the femininity of a created facial mesh.

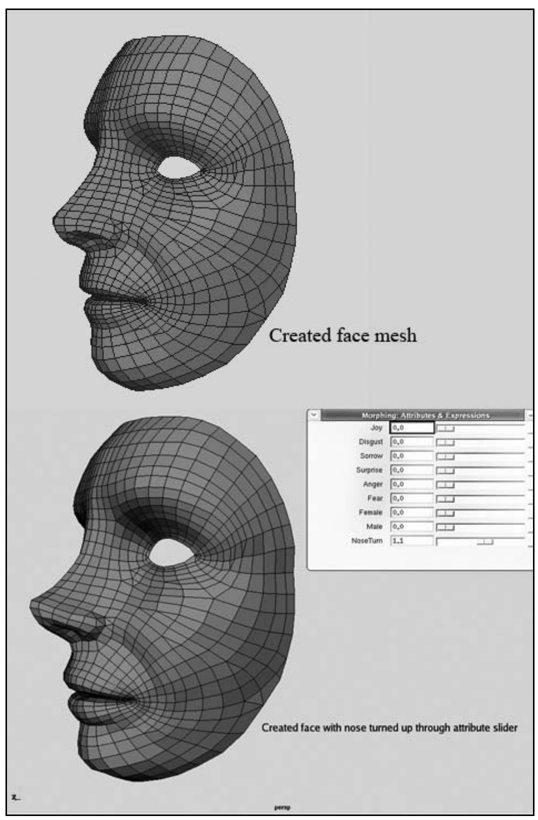

Figure 28: Turning up the nose of a created facial mesh using the user interface.

CHAPTER VI
RESULTS

Since the entire face data set contained many eigenvectors (63, or n-1 eigenvectors where n is the total number of exemplar faces), I had to choose how many components to retain (Fig. 29). The first few eigenvalues were much larger than the subsequent eigenvalues, (decreasing in value by a half or a fourth). The smaller eigenvalues were quite small and varied only by a few hundredths or thousandths. I chose to reduce the entire face vector space to only the first fifteen eigenvectors, or principal components. These were chosen based on how quickly the eigenvalues dropped towards zero.

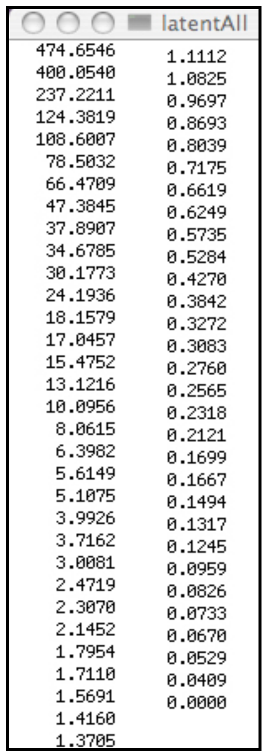
latentAll

474.6546	1.1112
400.0540	1.0825
237.2211	0.9697
124.3819	0.8693
108.6007	0.8039
78.5032	0.7175
66.4709	0.6619
47.3845	0.6249
37.8907	0.5735
34.6785	0.5284
30.1773	0.4270
24.1936	0.3842
18.1579	0.3272
17.0457	0.3083
15.4752	0.2760
13.1216	0.2565
10.0956	0.2318
8.0615	0.2121
6.3982	0.1699
5.6149	0.1667
5.1075	0.1494
3.9926	0.1317
3.7162	0.1245
3.0081	0.0959
2.4719	0.0826
2.3070	0.0733
2.1452	0.0670
1.7954	0.0529
1.7110	0.0409
1.5691	0.0000
1.4160	
1.3705	

Figure 29: Screen capture of the eigenvalues for the entire face data set.

More significant principal components could be scaled by much larger values than less significant components without making the created face mesh unrecognizable or unusable as a 'face'. The components with smaller eigenvalues are less likely to contain meaningful variance within the data set. Largely influencing a face by a less significant component increases the chance that the new face will fall outside the range of form we find acceptable as a 'face' (Fig. 30). Unacceptable faces would also include any meshes containing overlapping or severely distorted geometry. These meshes could not be rendered correctly by the software.

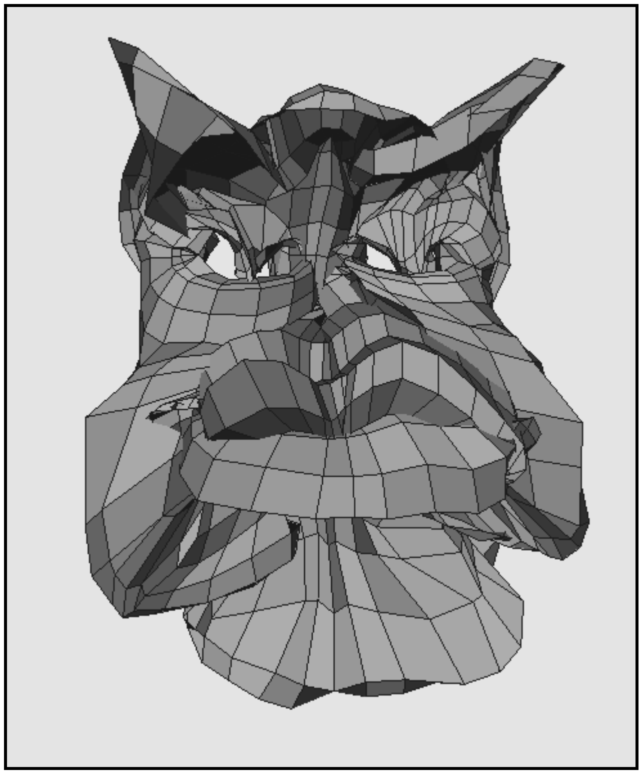

Figure 30: Face created using a large value for the coefficient of a less significant component.

Initially, this result was avoided by not allowing the coefficient value to be larger or smaller than specific positive or negative bounds. These bounds were chosen based on my own observation as to which coefficient values for each component produced a mesh which was either unrecognizable as a face or could not be handled correctly by the software. In the case of the entire face data set, the first principal component was limited to coefficient values between -70.0 and 100.0, while the fifteenth principal component was limited in value by being scaled between -25.0 and 30.0. This method nearly always limits the user to creating a 'safe' face.

An alternate method of clamping the ranges of the coefficient values was to normalize the values based on the square root of the component's eigenvalue. Since the meaningful variation of the data associated with each component is assumed to fit a Gaussian distribution, the normalized coefficient values directly corresponded to the number of standard deviations from the mean. For each component, a normalized coefficient range from -1.0 to 1.0 would correspond statistically to the majority (68.2%) of the variation associated with that component. Normalized coefficient values that lie outside the -3.0 to 3.0 range produced results more than three standard deviations from the mean. Almost all (99%) of the data's variation lies within three standard deviations of the mean. Normalizing the coefficient values creates a statistically meaningful range for all components.

Observed Influence of Principal Components – Entire Face Data Set

Principal component analysis is a statistical analysis that identifies general variation of a data set. Attributes associated with artistic style and those associated with facial expression are mixed throughout the data set and are not easily separated and identified by principal component analysis. Despite this, general properties of each principal component of the entire created face data set can be identified. Although patterns could be identified, there was an accompanying variance in artistic style of the face that is harder to quantify. It is possible to adjust the coefficient values of certain components to create a face resembling the exemplar sculptures by a specific artist.

In the following discussion, all illustrations of the variation due to modifying a principal component coefficient have the face with the negatively scaled component on the left, and the face with the positively scaled component on the right. The coefficient factors shown are normalized and specified in terms of standard deviation values. There are figures for each principal component. One figure shows a range of standard deviation chosen based on the appearance of the resulting faces. The second figure for each principal component shows a range of +/- 3 standard deviations. Illustrations show faces with uniform scale within each principal component set.

First and Second Principal Components

The first and second principal components correspond to the face data set's greatest amount of variance. These components were identified as directly varying the width and height of the face. The first component varies the vertical asymmetry (ratio of forehead width to chin width) and overall size of the face mesh, while the second component varies the aspect ratio of the mesh (Fig. 31-34).

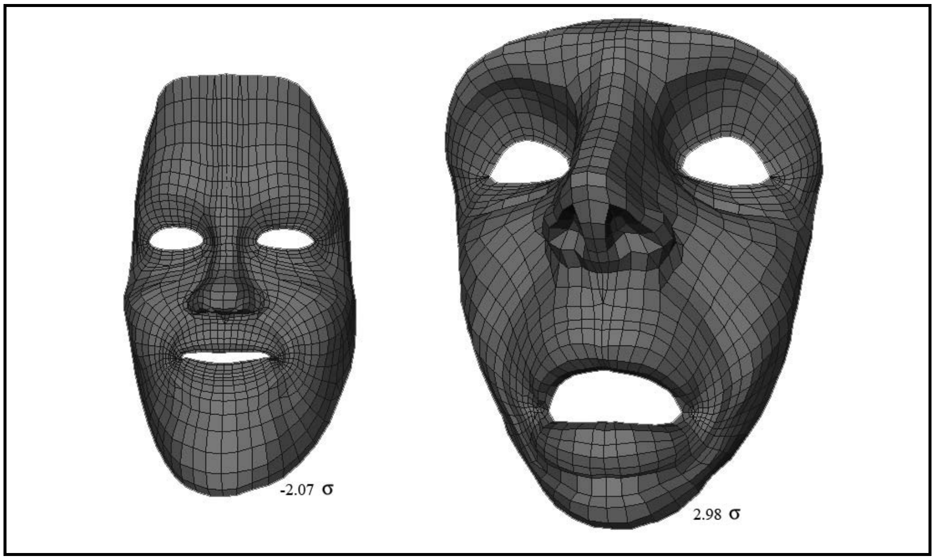

Figure 31: Varying the first principal component coefficient.

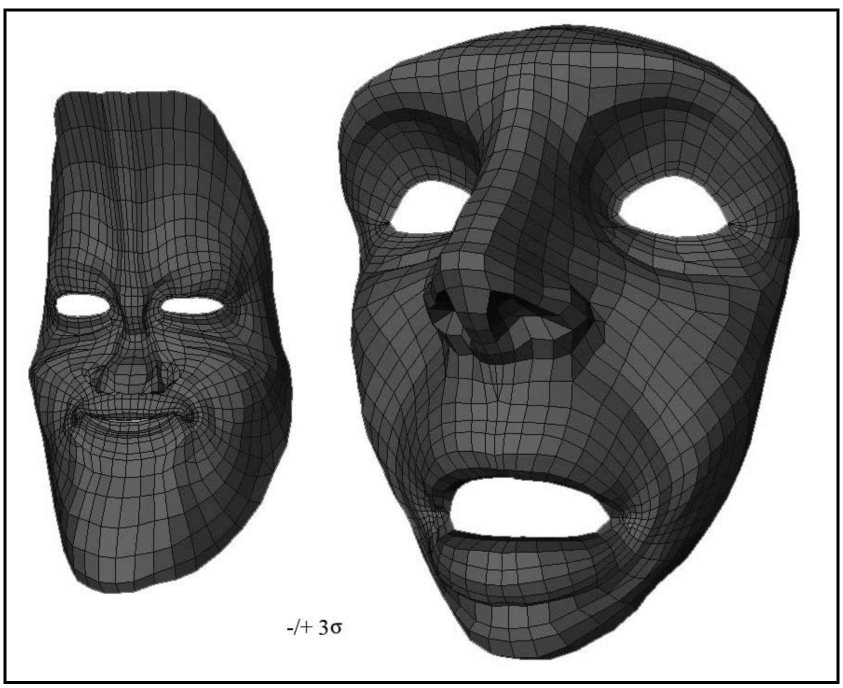

Figure 32: Varying the first principal component coefficient by -/+ 3 standard deviations.

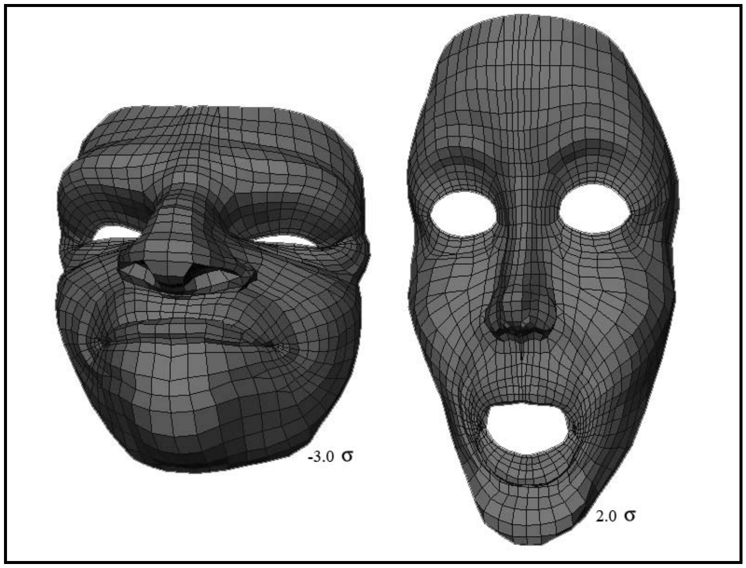

Figure 33: Varying the second principal component coefficient.

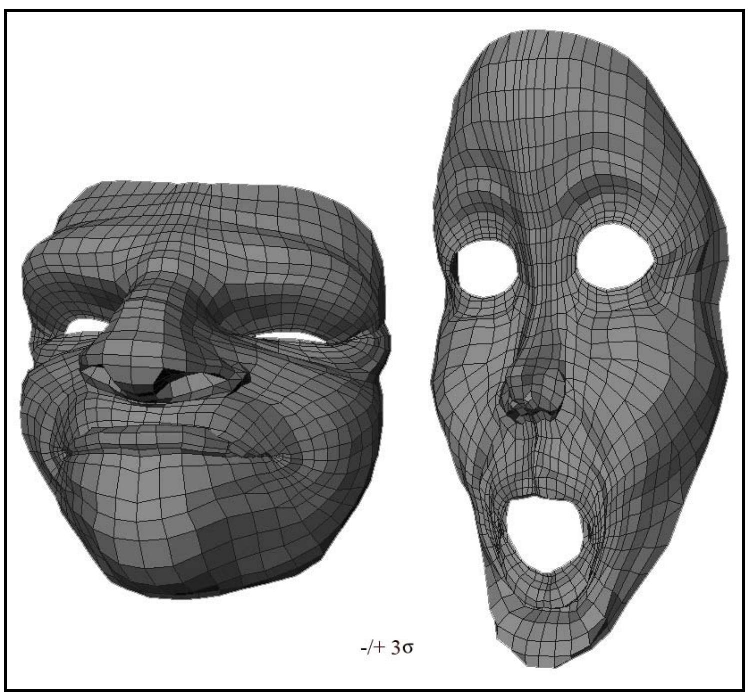

Figure 34: Varying the second principal component coefficient by -/+ 3 standard deviations.

Third Principal Component

The third principal component also varies the aspect ratio of the face, but the aspect ratio is changed inversely from varying second component coefficient (Fig. 35-36). Both the second and third components also vary a specific facial feature; the mouth. In addition to changing the overall aspect ratio of the face, it also varies how open or closed the mouth is.

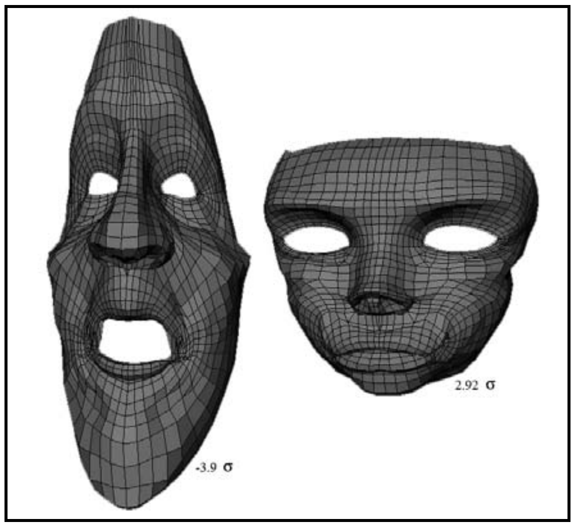

Figure 35: Varying the third principal component coefficient.

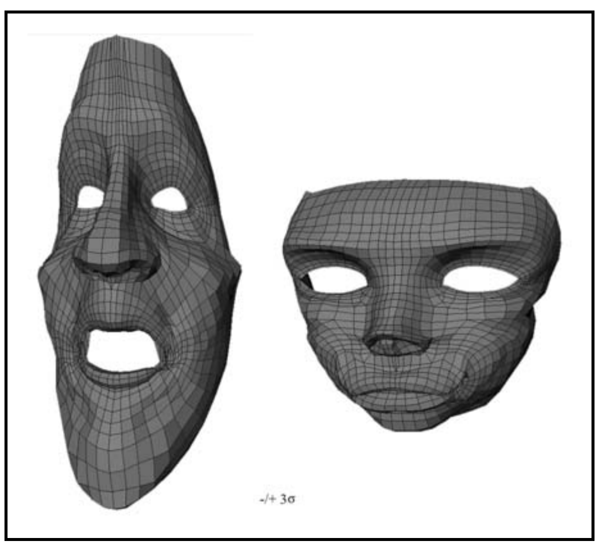

Figure 36: Varying the third principal component coefficient by -/+ 3 standard deviations.

Fourth Principal Component

The fourth component corresponds to an obvious change in facial expression. There was a variation toward fear in the negative direction and a variation toward disgust in the positive direction (Fig. 37-38). There was an accompaniment of change in artistic style of the face as well as a drastic change in the nose from up-turned to hooked.

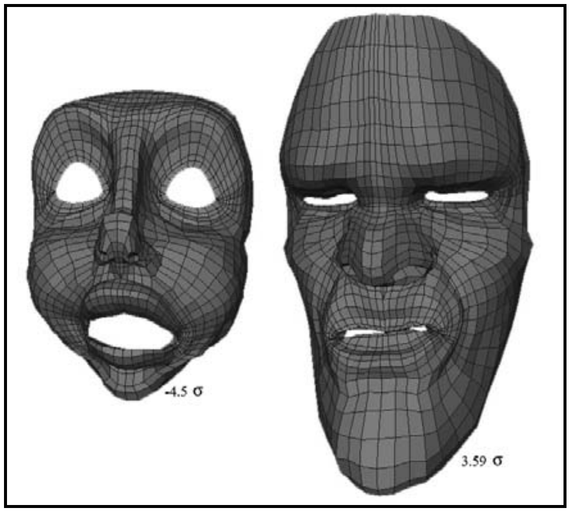

Figure 37: Varying the fourth principal component coefficient.

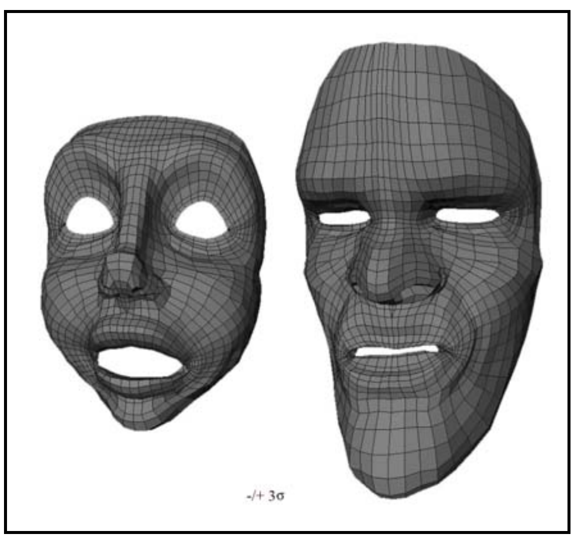

Figure 38: Varying the fourth principal component coefficient by -/+ 3 standard deviations.

Fifth Principal Component

Similarly, there was a variance in the joy and sorrow of the facial mesh when scaling the fifth principal component's coefficient (Fig. 39-40). There was also a change in the cheekbones that affected the entire mesh. The width of the cheekbones varied the face from a more square shape to more of a diamond.

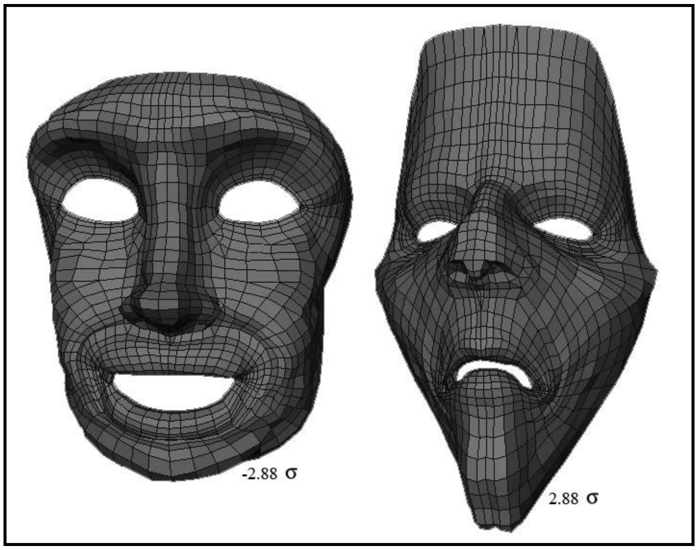

Figure 39: Varying the fifth principal component coefficient.

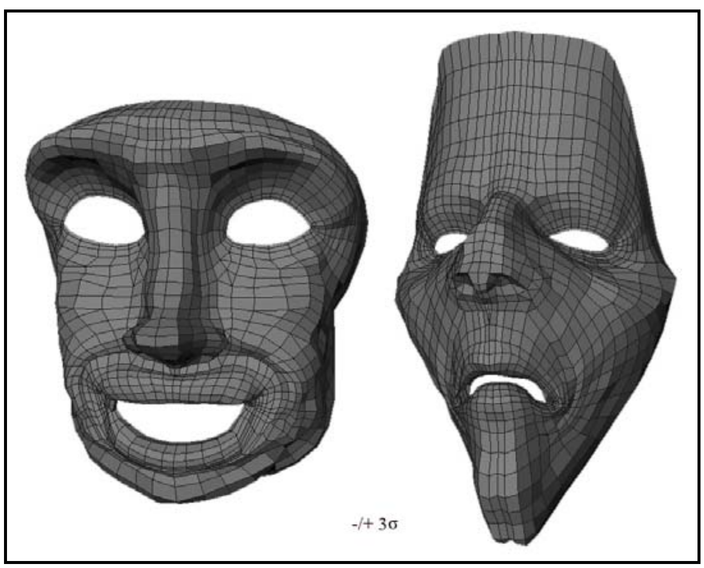

Figure 40: Varying the fifth principal component coefficient by -/+ 3 standard deviations.

Sixth Principal Component

The sixth component corresponds to the open and closed position of the mouth, along with another variance in artistic style (Fig. 41-42).

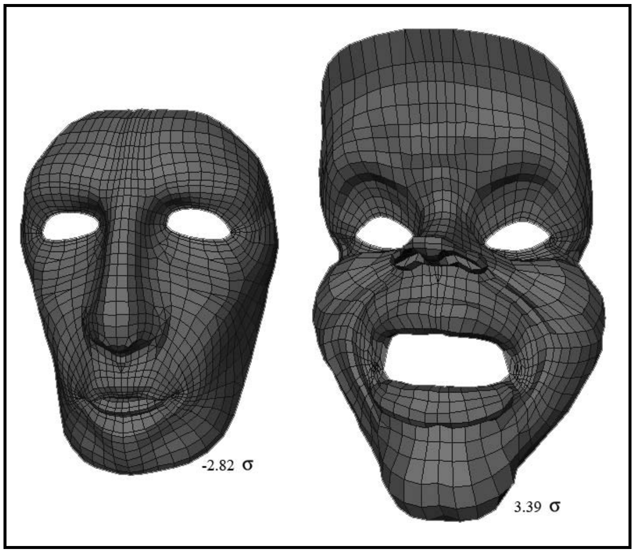

Figure 41: Varying the sixth principal component coefficient.

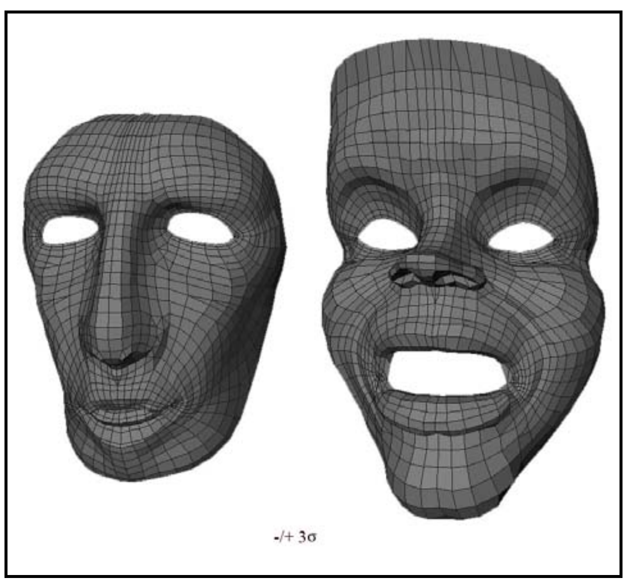

Figure 42: Varying the sixth principal component coefficient by -/+ 3 standard deviations.

Seventh Principal Component

The seventh principal component has the affect of expanding and compressing the facial features, both in the size of the features and in their protrusion from the face (Fig. 43-44).

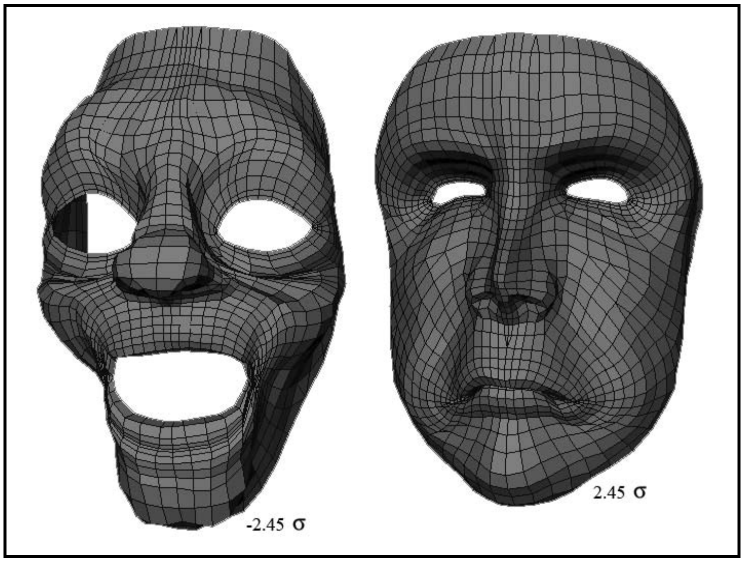

Figure 43: Varying the seventh principal component coefficient.

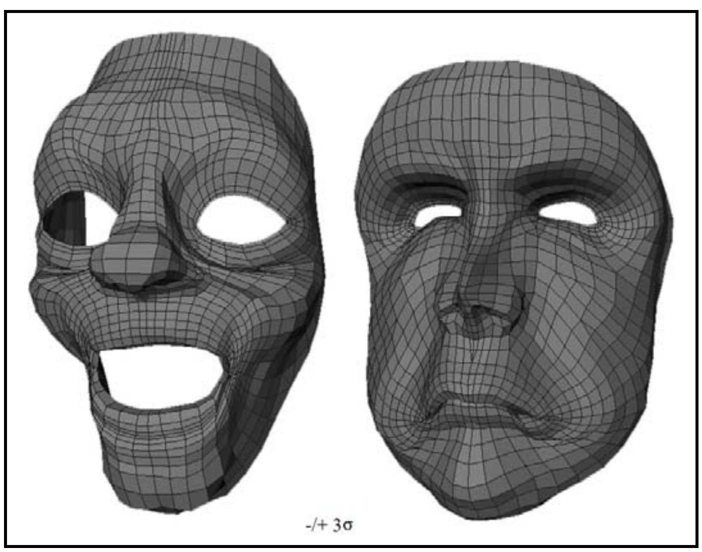

Figure 44: Varying the seventh principal component coefficient by -/+ 3 standard deviations.

Eighth Principal Component

Varying the eighth principal component resulted in a major effect on the profile of the face (Fig. 45-46). A negative coefficient value produces a concave profile of the face, while a positive coefficient produces a convex result. A corresponding change in the direction of the turn of the nose is also shown.

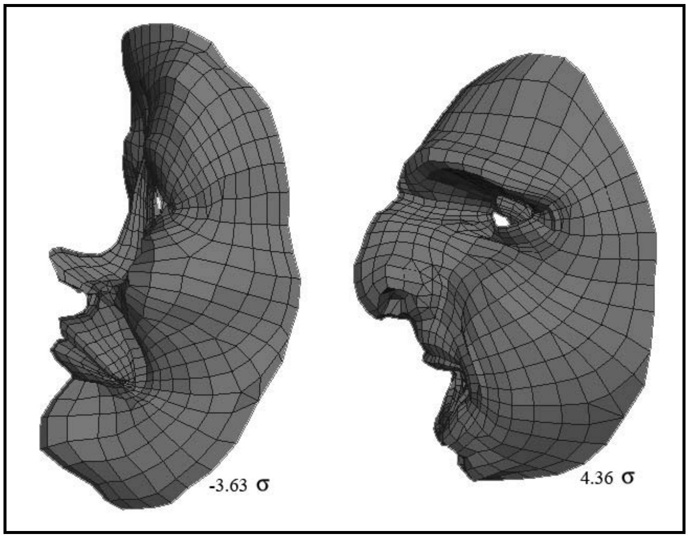

Figure 45: Varying the eighth principal component coefficient.

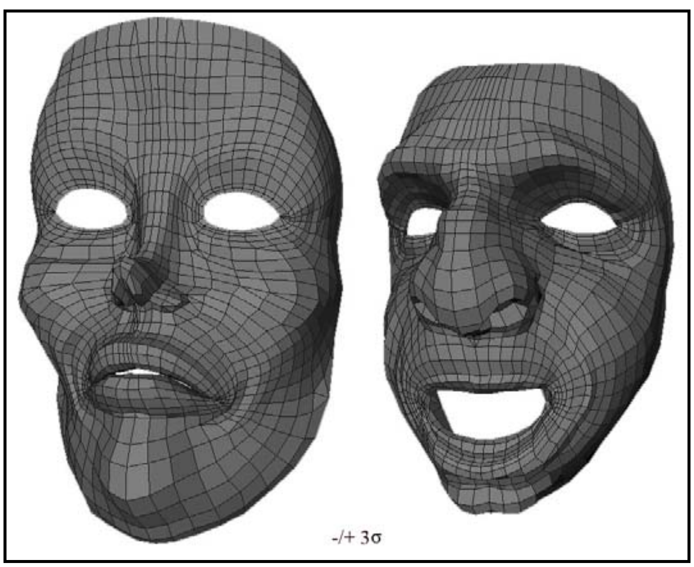

Figure 46: Varying the eighth principal component coefficient by -/+ 3 standard deviations.

Ninth Principal Component

The ninth principal component shows a distinct variation between sorrow and joy, as well as the introduction of distinct facial wrinkles (Fig. 47-48).

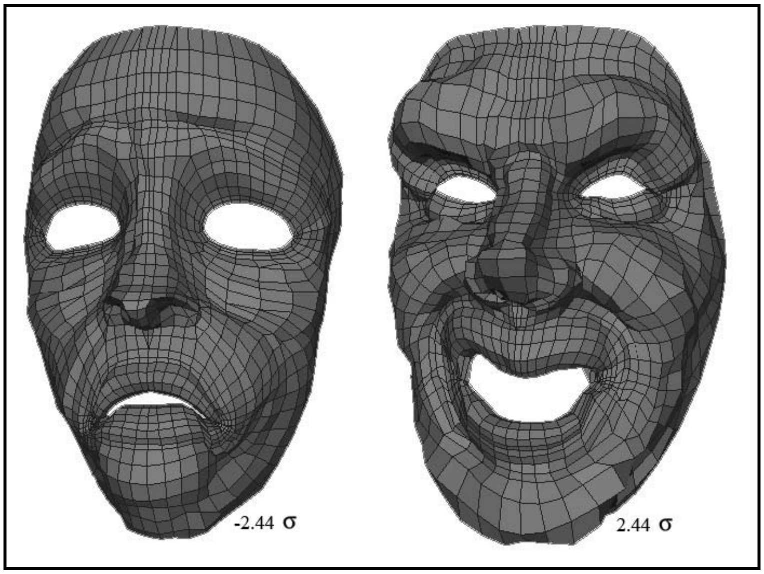

Figure 47: Varying the ninth principal component coefficient.

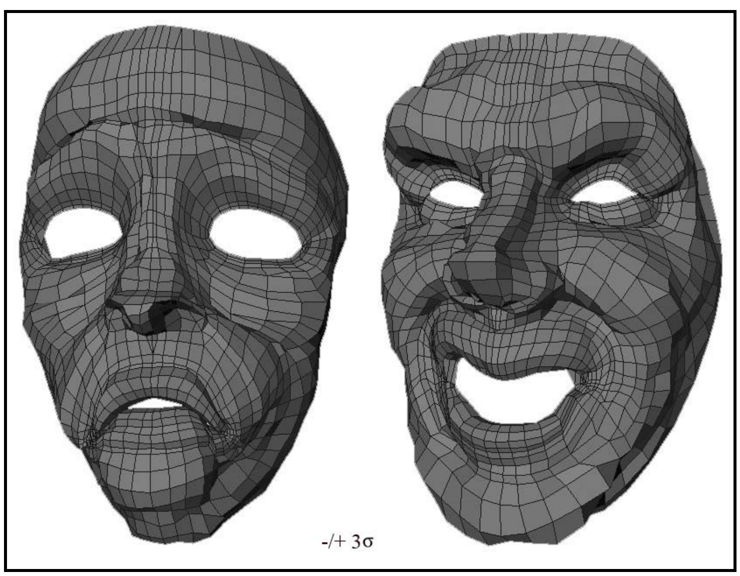

Figure 48: Varying the ninth principal component coefficient by -/+ 3 standard deviation.

Tenth Principal Component

The tenth principal component produces a twisting of the features (as seen within the disgust expression exemplar faces) for negative coefficient values, as well as opening and closing the mouth (Fig. 49-50).

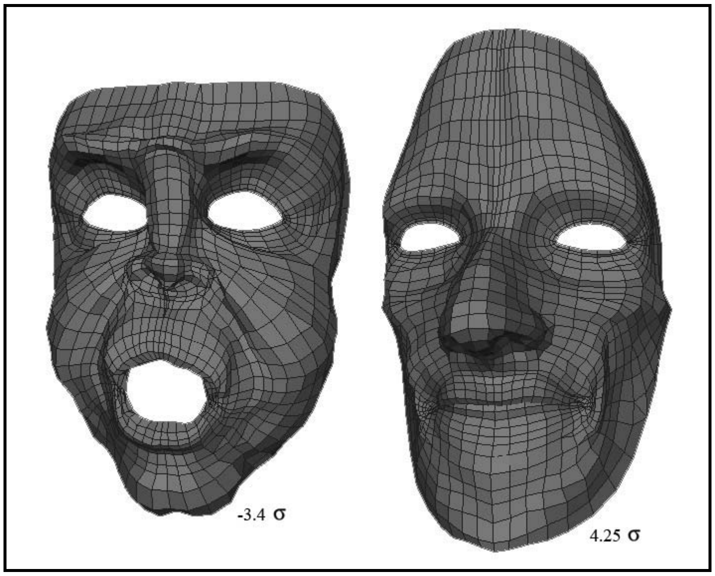

Figure 49: Varying the tenth principal component coefficient.

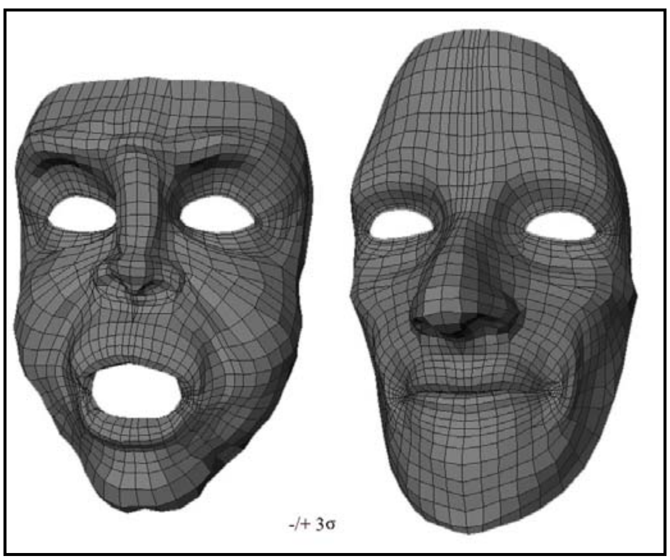

Figure 50: Varying the tenth principal component coefficient by -/+ 3 standard deviations.

Eleventh Principal Component

The result of scaling the eleventh component combines a mixture of indications of fear and joy with a mix of anger and disgust expression (Fig. 51-52). The effect of this is difficult to quantify as a specific attribute.

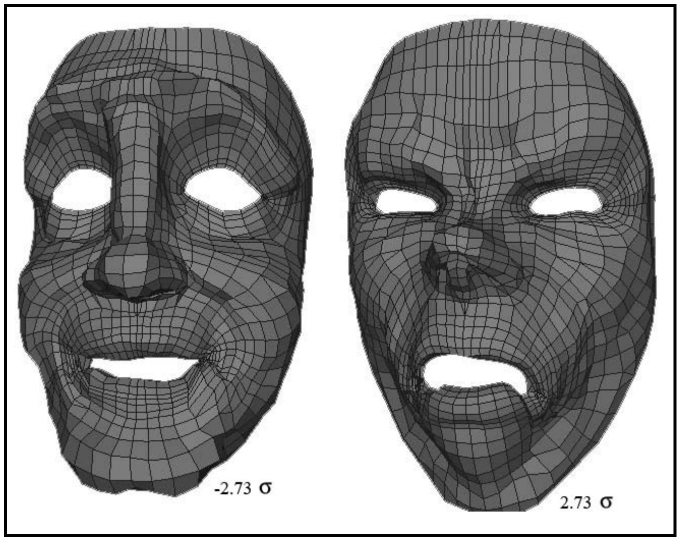

Figure 51: Varying the eleventh principal component coefficient.

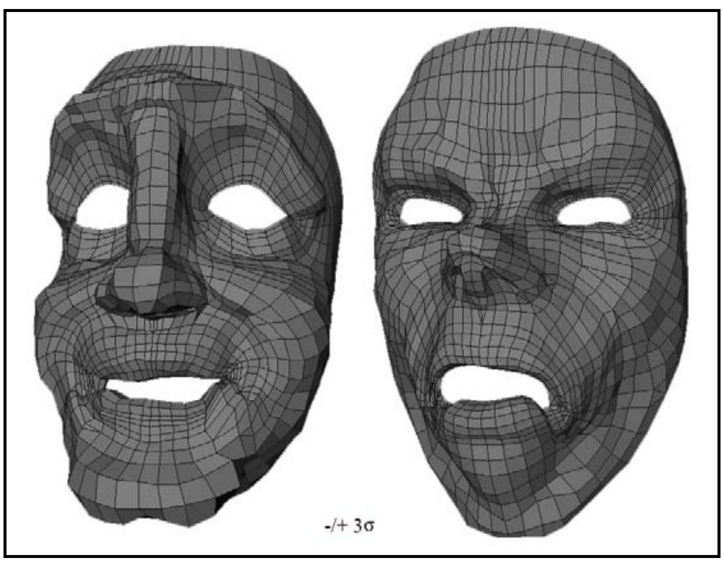

Figure 52: Varying the eleventh principal component coefficient by -/+ 3 standard deviations.

Twelfth Principal Component

The twelfth principal component shows an interesting variance in the profile of the facial mesh. Scaling the component negatively produces very short nose, mouth, and chin features which are very close together. Scaling in the positive direction elongates these features, moving them farther apart and protruding them all further from the rest of the face (Fig. 53-54).

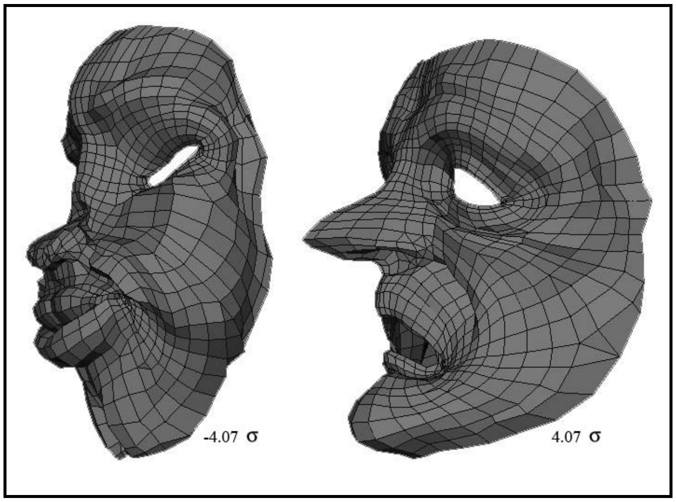

Figure 53: Varying the twelfth principal component coefficient.

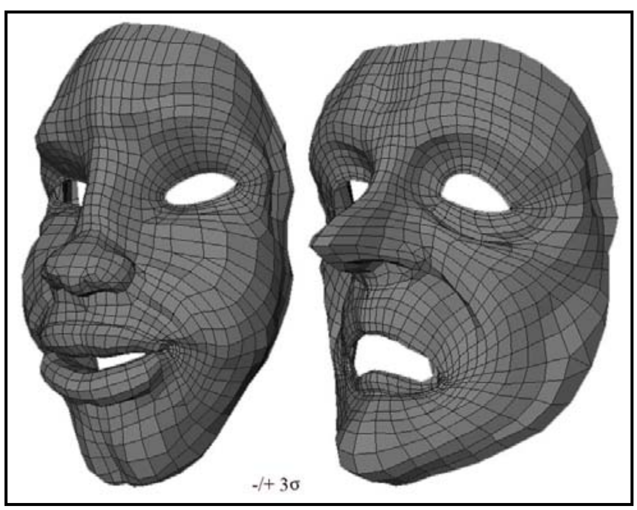

Figure 54: Varying the twelfth principal component coefficient by -/+ 3 standard deviations.

Thirteenth Component

Varying the thirteenth component resulted in a variation from joy to anger as well as a general squashing and expanding of the face (Fig. 55-56).

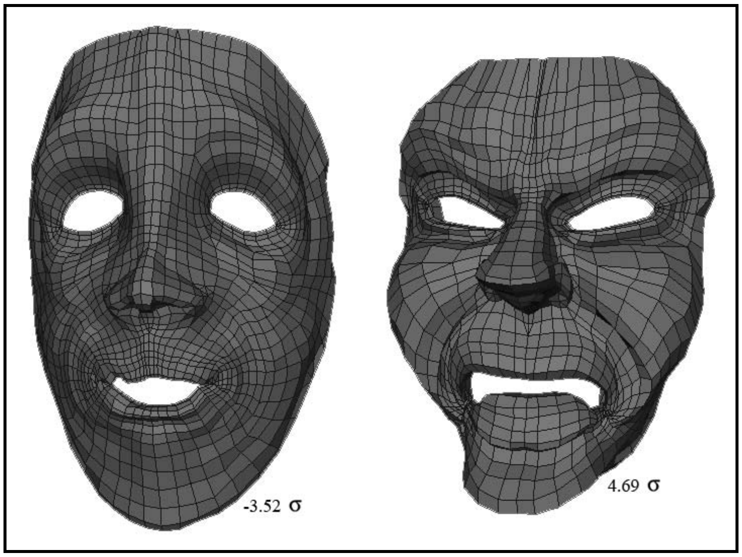

Figure 55: Varying the thirteenth principal component coefficient.

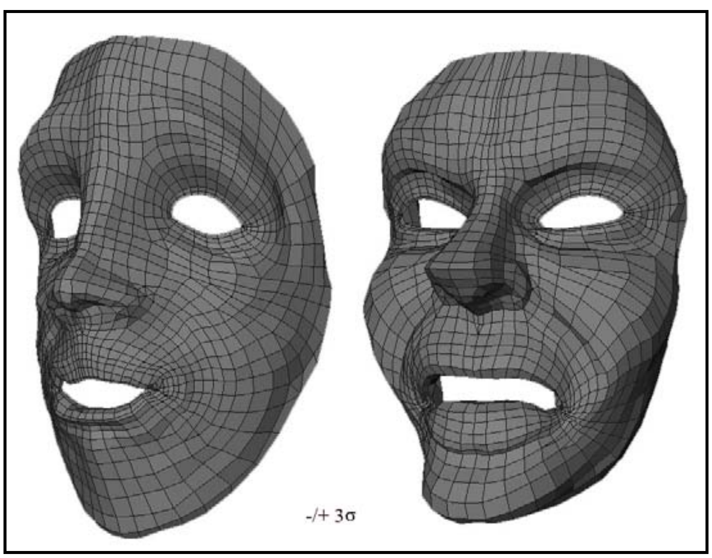

Figure 56: Varying the thirteenth principal component coefficient by -/+ 3 standard deviations.

Fourteenth and Fifteenth Principal Components

The fourteenth and fifteenth principal components are the least significant components used. They correlate to changes in profile and are the most sensitive. The fourteenth component varies the area and amount of protrusion of the forehead and chin features relative to the area and protrusion of the nose and lips (Fig. 57-58). The fifteenth component increases and decreases the protrusion of all features of the mesh along with some change in expression (Fig. 59-60).

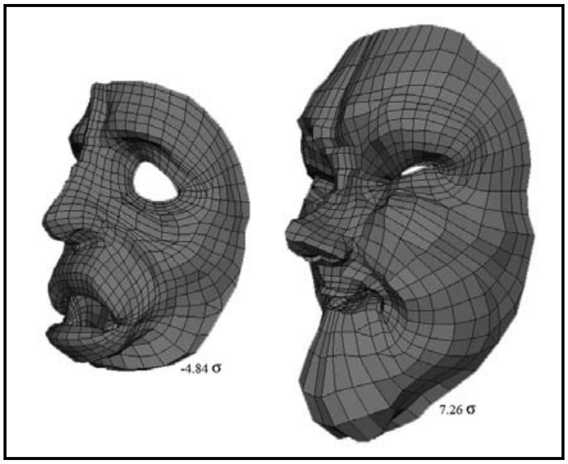

Figure 57: Varying the fourteenth principal component coefficient.

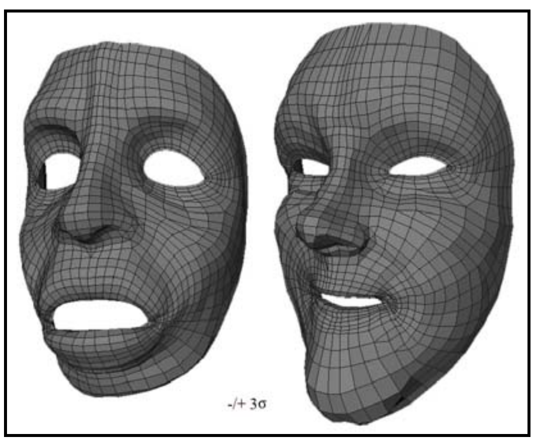

Figure 58: Varying the fourteenth principal component coefficient by -/+ 3 standard deviations.

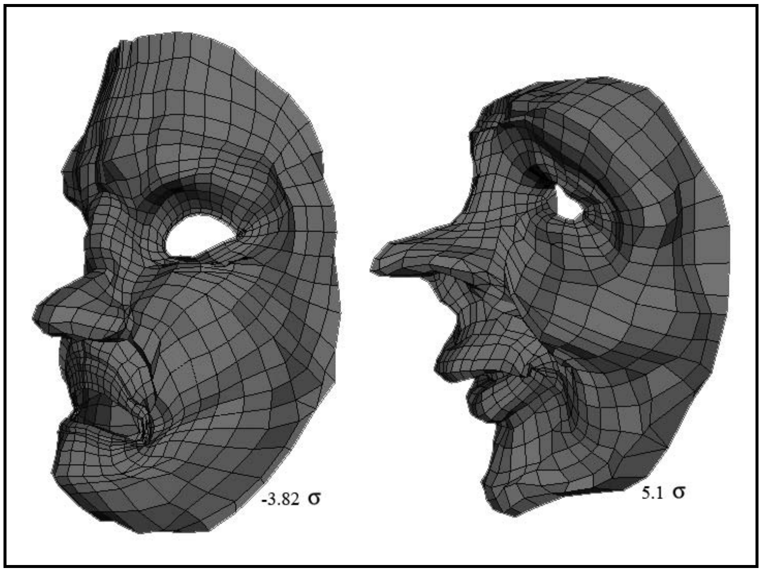

Figure 59: Varying the fifteenth principal component coefficient.

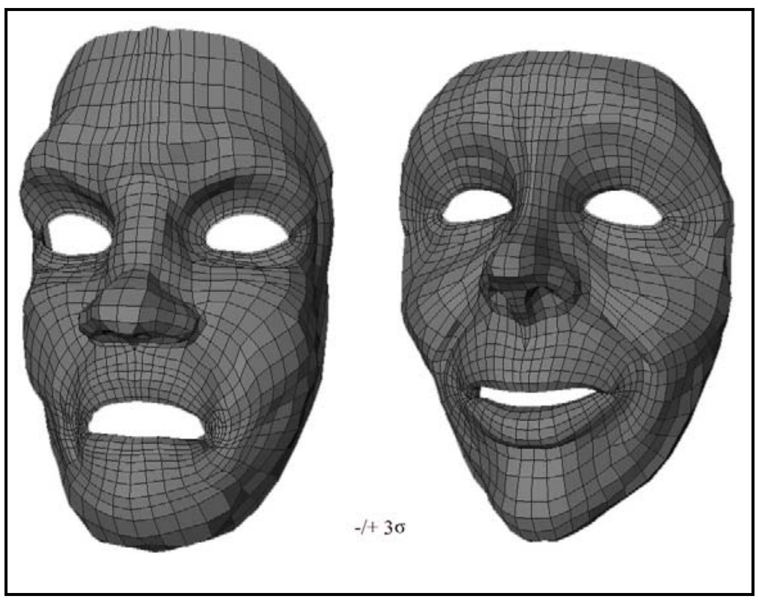

Figure 60: Varying the fifteenth principal component coefficient by -/+ 3 standard deviations.

Observed Influence of Principal Components – Neutral Face Space

The face data expression sub-sets of the entire created face data set had a smaller number of eigenvectors. The vector spaces representing these sub-sets of data were defined using all their respective eigenvectors. Figure 61 shows the eigenvalues for the neutral expression face data sub-set. Because the sub-sets had so few principal components, I chose to retain all of them.

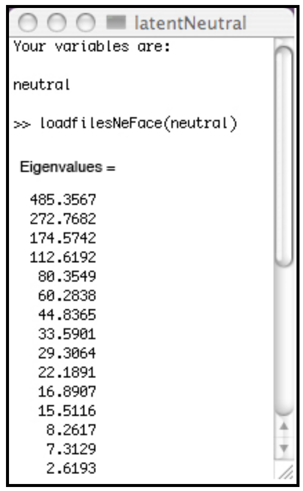

Figure 61: Screen capture of the eigenvalues for the neutral face space.

Exploration of the principal components of only the neutral expression faces excludes variation due to artist created expression. However, the faces produced from this face space still sometimes appear to contain an expression. This is likely due to a variation in the position of features that psychologically imply expression. For instance, if a scaling a component results in corners of the mouth that are higher on the face while the middle of the mouth stays fixed, this could appear as a measure of joy. Faces resulting from scaling the components of the neutral expression data allowed the identification of change in artistic style, although most components of the space did not directly correspond to a specific artist. Scaling the first three components produced similar results to the exploration of the entire face data set. The components corresponded with vertical asymmetry and aspect ratio, but contained no variation due to expression, and slightly different variation due to artistic style. Other components showed variation of profile and facial feature size, proportion, and style just as they had in the entire face data set. The biggest difference was the absence of artist intended facial expression.

First and Second Principal Components

Like the entire created face set, the first and second principal component of neutral face set corresponded to vertical asymmetry and aspect ratio. For the neutral face set, a negatively scaled first component resulted in a large face with wider features at the top of the mesh. A positively scaled first component created a smaller, narrower face with wider features at the base of the mesh (Fig. 62-63). The correspondence of the second principal component was inverted from that of the entire face set. A negatively scaled second component created a longer, narrower face while a positive coefficient resulted in a shorter, wider face (Fig. 64-65). Each of these components affects the openness of the mouth, but not as dramatically as with the entire face data set. This is likely due to the absence of facial expression. In the entire data set the surprise and fear expressions probably contribute heavily to the variation towards a wide open mouth.

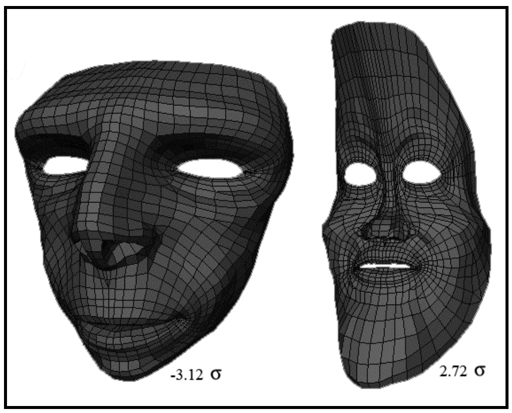

Figure 62: Varying the first principal component coefficient of the neutral face set.

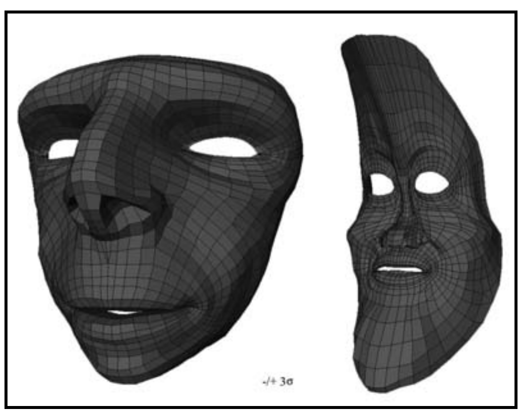

Figure 63: Varying the first principal component coefficient of the neutral face set by -/+ 3 standard deviations.

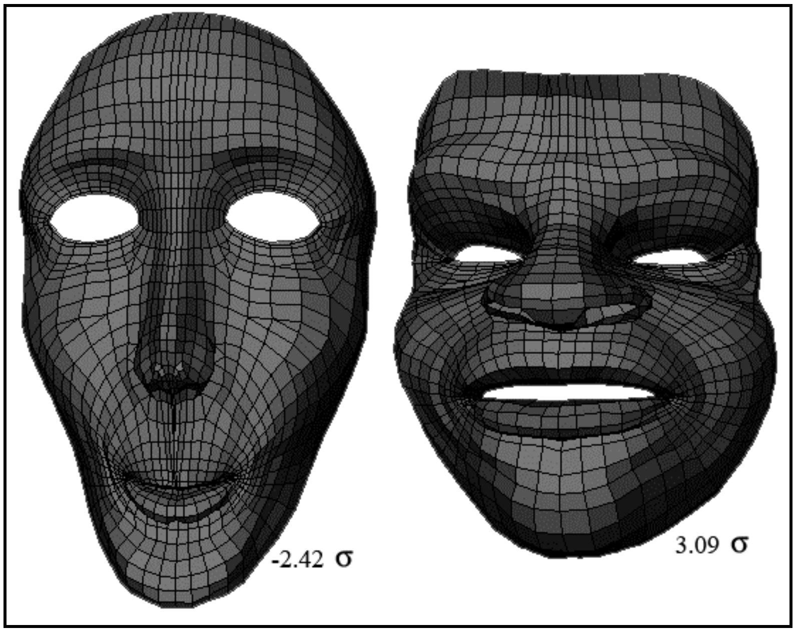

Figure 64: Varying the second principal component coefficient of the neutral face set.

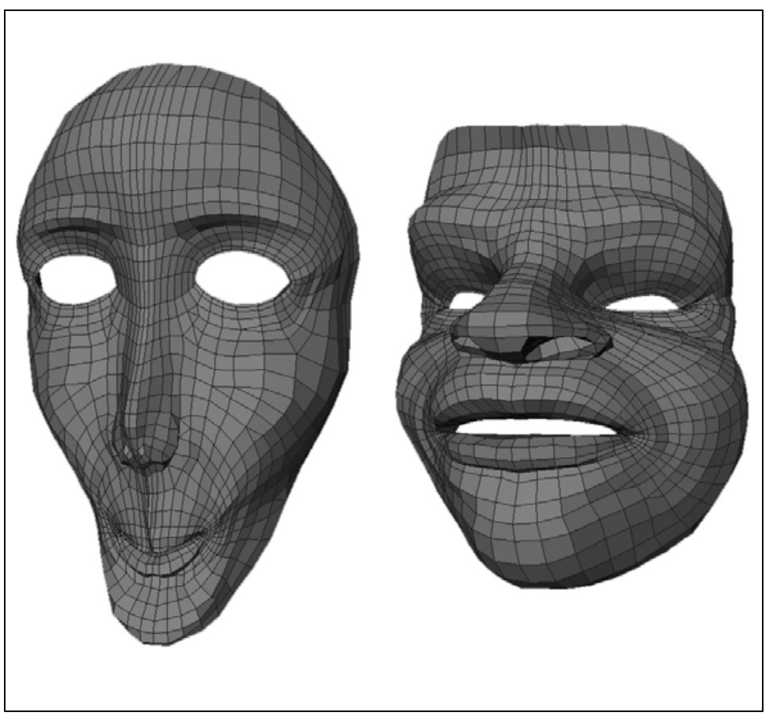

Figure 65: Varying the second principal component coefficient of the neutral face set by -/+ 3 standard deviations.

Third Principal Component

Interestingly, varying the third principal component of the neutral face set resulted in nearly identical faces as with varying the third principal component of the entire face set. In the case of the neutral set, the mouth does not open as widely when the component is scaled negatively (Fig. 66-67). There is variation of aspect ratio that is similar to the results seen by varying the second component of the neutral data set.

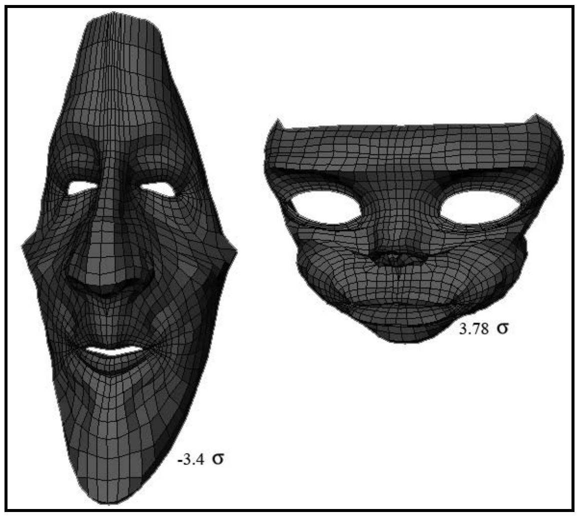

Figure 66: Varying the third principal component coefficient of the neutral face set.

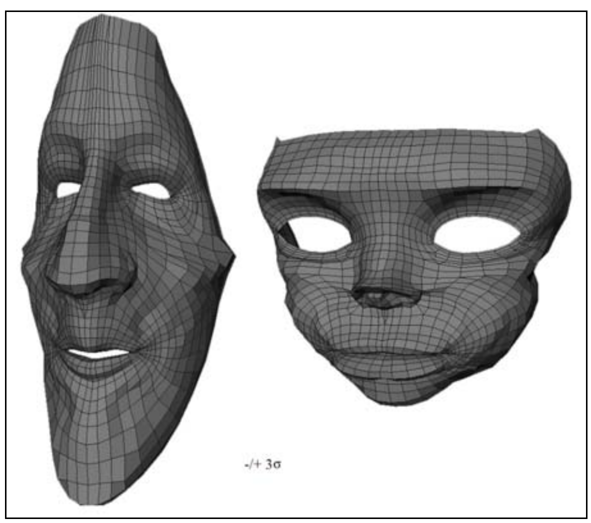

Figure 67: Varying the third principal component coefficient of the neutral face set by -/+ 3 standard deviations.

Fourth Principal Component

The fourth component varies the width of the face, the size of the forehead and chin areas, and the direction of the turn of the nose (either upwards toward the forehead or downwards toward the chin) (Fig. 68-69). The faces created by scaling the fourth component of the neutral face set are almost the same faces created by scaling the fourth component of the entire face. The neutral faces had the opposite turn of the noses as seen from varying the coefficient of the fourth component of the entire face set.

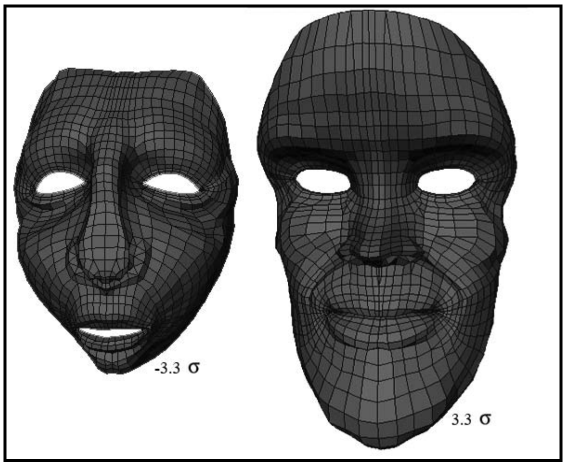

Figure 68: Varying the fourth principal component coefficient of the neutral face set.

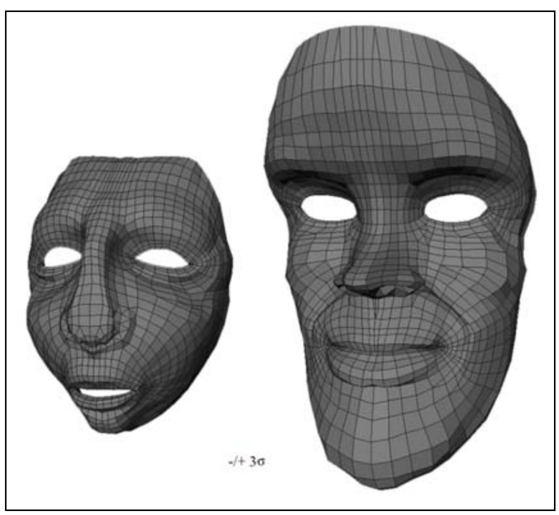

Figure 69: Varying the fourth principal component coefficient of the neutral face set by -/+ 3 standard deviations.

Fifth Principal Component

Varying the fifth principal component affects the overall shape of the face mesh (Fig. 70-71). A negative coefficient created a shorter, square, face while a positive coefficient resulted in a longer, diamond-like face. The vertical length of the nose was also affected. The faces created by varying the fifth component are expressionless versions of the faces created by varying the fifth component of the entire face set.

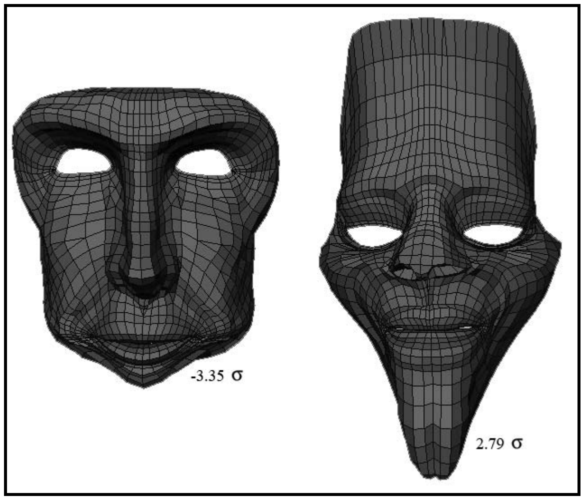

Figure 70: Varying the fifth principal component coefficient of the neutral face set.

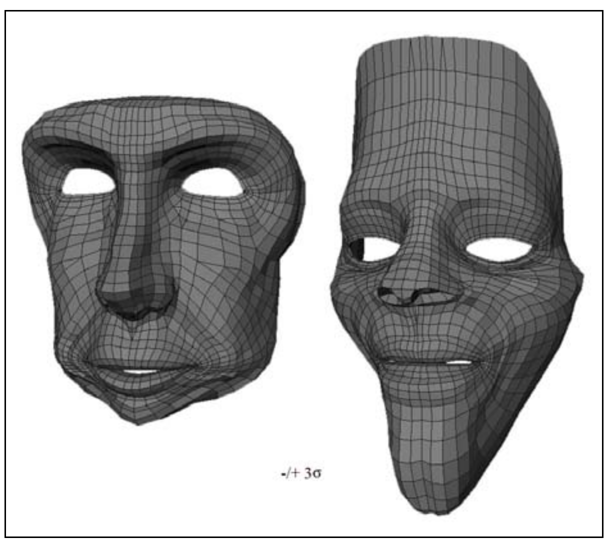

Figure 71: Varying the fifth principal component coefficient of the neutral face set by -/+ 3 standard deviations.

Sixth Principal Component

The sixth principal component of the neutral face set does not follow the pattern of similarity to the components of the entire face data set. Here it resulted in an inverse relationship between nose and lip size. Negative and positive coefficients altered the direction of turn of the nose. The face created by positively scaling the sixth component is more female in gender (Fig. 72-73).

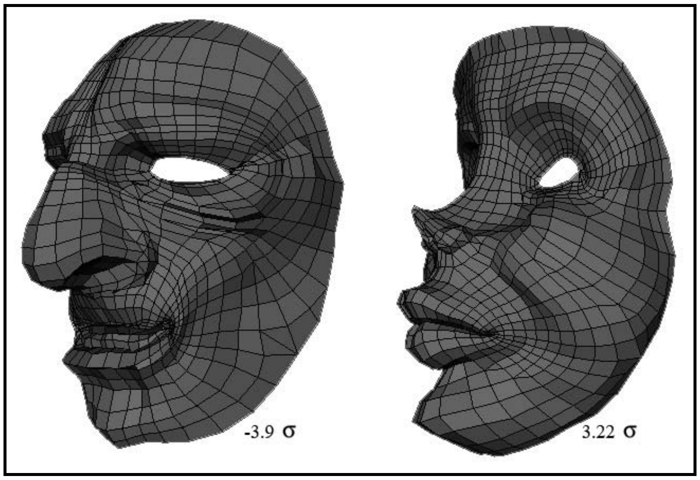

Figure 72: Varying the sixth principal component coefficient of the neutral face set.

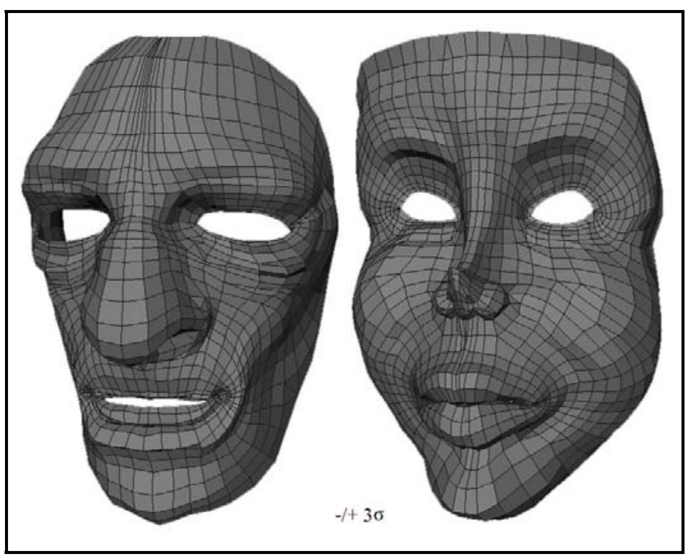

Figure 73: Varying the sixth principal component coefficient of the neutral face set by -/+ 3 standard deviations.

Seventh Principal Component

The seventh principal component seems to be the inverse of the sixth component but with slightly less pronounced features and less change of gender (Fig. 74-75). The variation of the seventh component reflects different artistic styles than the variation of the sixth component.

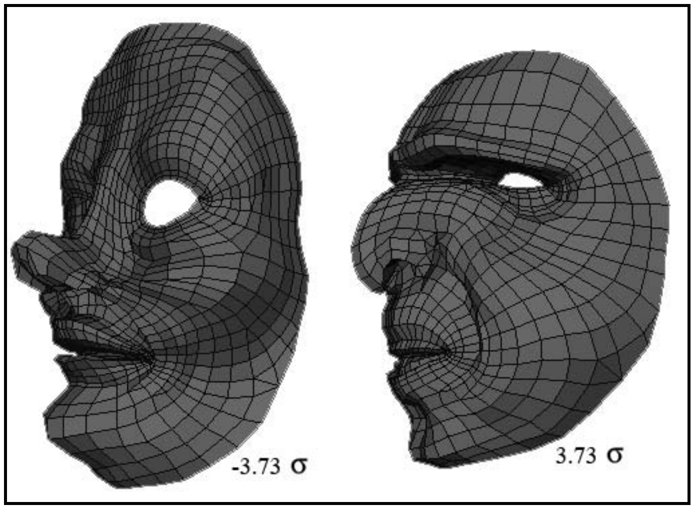

Figure 74: Varying the seventh principal component coefficient of the neutral face set.

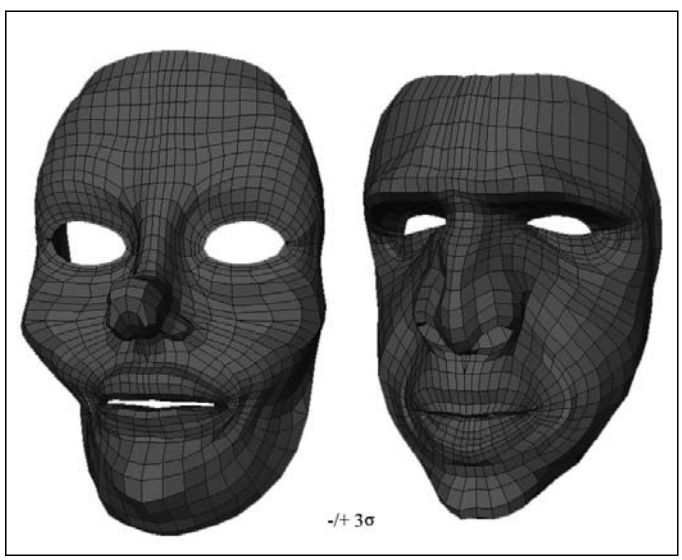

Figure 75: Varying the seventh principal component coefficient of the neutral face set by -/+ 3 standard deviations.

Eighth Principal Component

The eighth principal component of the neutral set corresponds to the protrusion of facial features (Fig. 76-77). A negative coefficient created thinner, more sunken features while a positive coefficient created a fleshier look. There is also a variation from a wider back of the facial mesh to a narrower one.

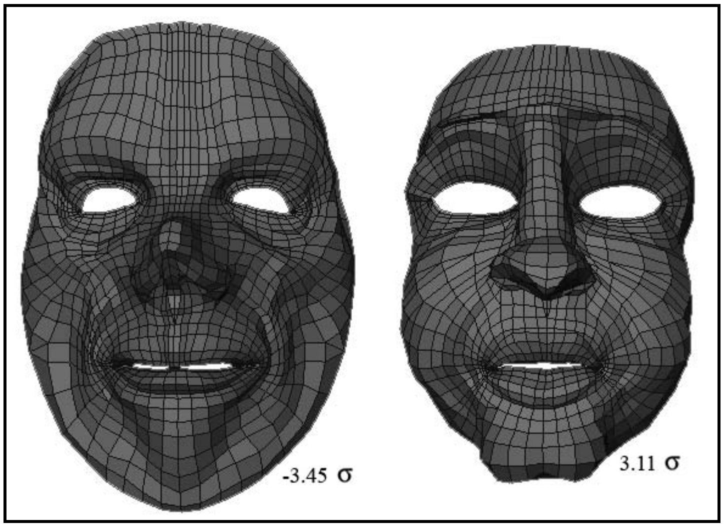

Figure 76: Varying the eighth principal component coefficient of the neutral face set.

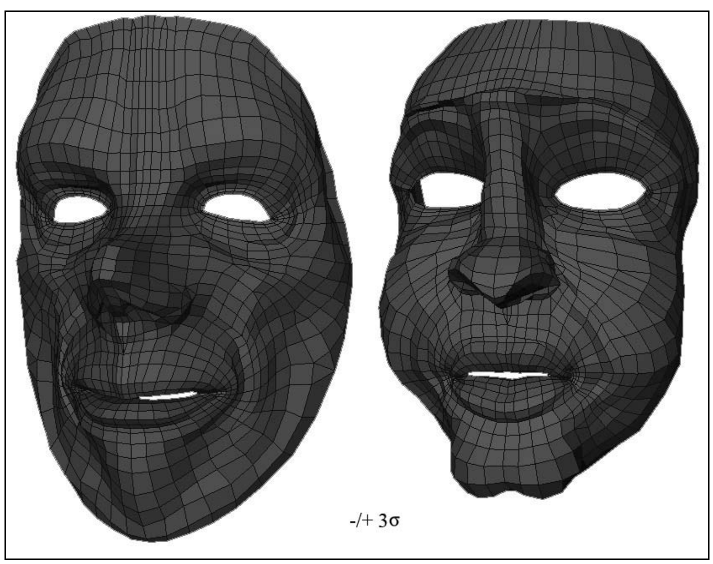

Figure 77: Varying the eighth principal component coefficient of the neutral face set by -/+ 3 standard deviations.

Ninth Principal Component

The ninth principal component corresponds to a change of overall shape of the facial mesh (Fig. 78-79). A negative coefficient resulted in a more oval shape with wider, larger facial features, while a positive coefficient created a square shape with thinner, smaller facial features.

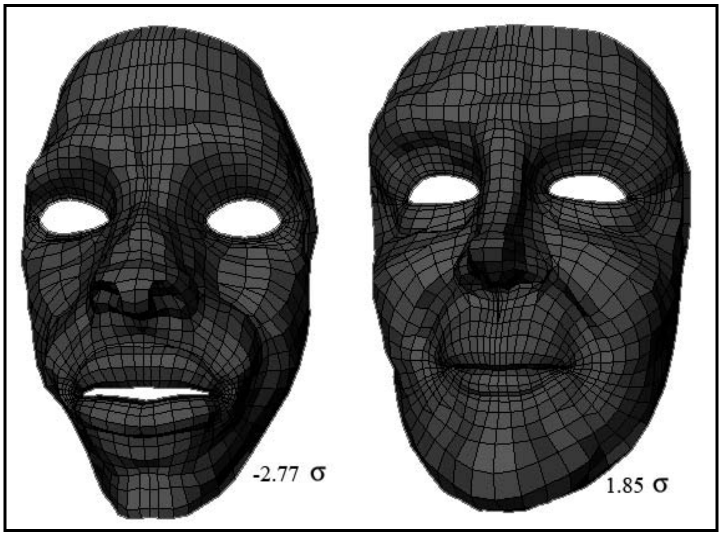

Figure 78: Varying the ninth principal component coefficient of the neutral face set.

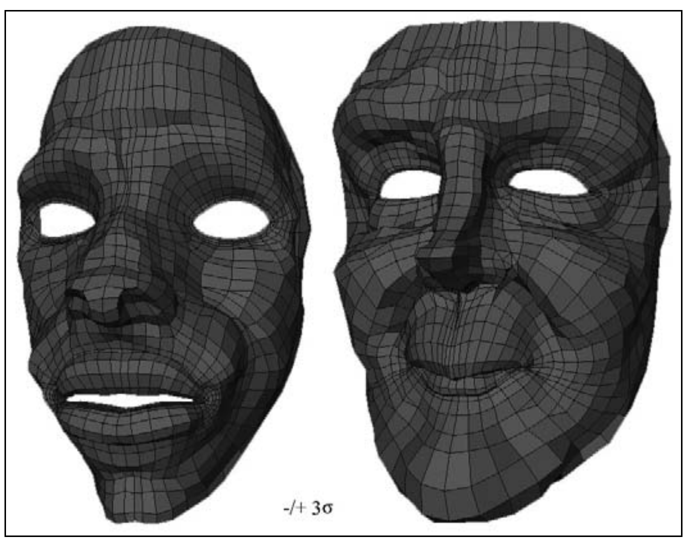

Figure 79: Varying the ninth principal component coefficient of the neutral face set by -/+ 3 standard deviations.

Tenth Principal Component

Varying the tenth principal component resulted in a variation of the protrusion of the facial features (Fig. 80-81). A negative coefficient protruded the nose and chin much further than the forehead and lips. A positive coefficient protruded the lips and forehead much more than the nose and chin. A similar relationship was seen when varying the twelfth principal component of the entire face data set.

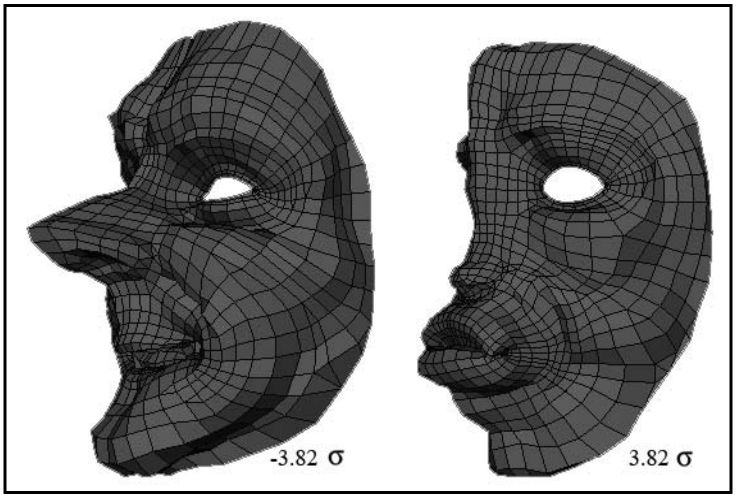

Figure 80: Varying the tenth principal component coefficient of the neutral face set.

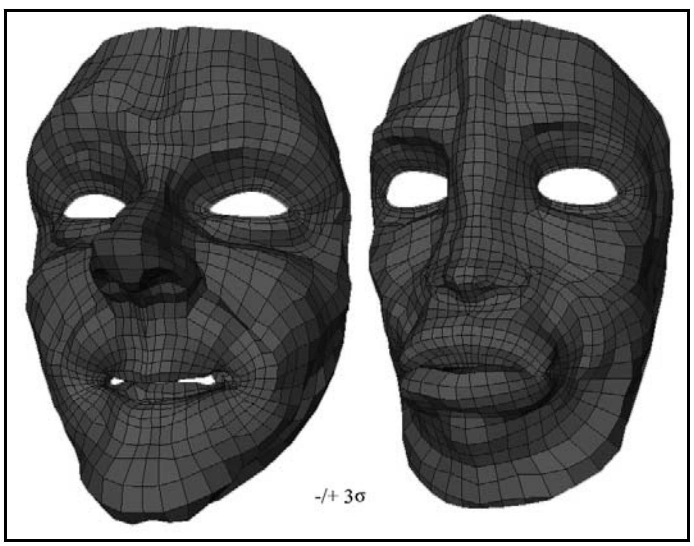

Figure 81: Varying the tenth principal component coefficient of the neutral face set by -/+ 3 standard deviations.

Eleventh Principal Component

Varying the eleventh component caused a variation toward larger more pronounced facial features overall as well as a change in artistic style (Fig. 82-83). It also appears to be the inverse of the eleventh principal component of the entire face data set minus the expressions.

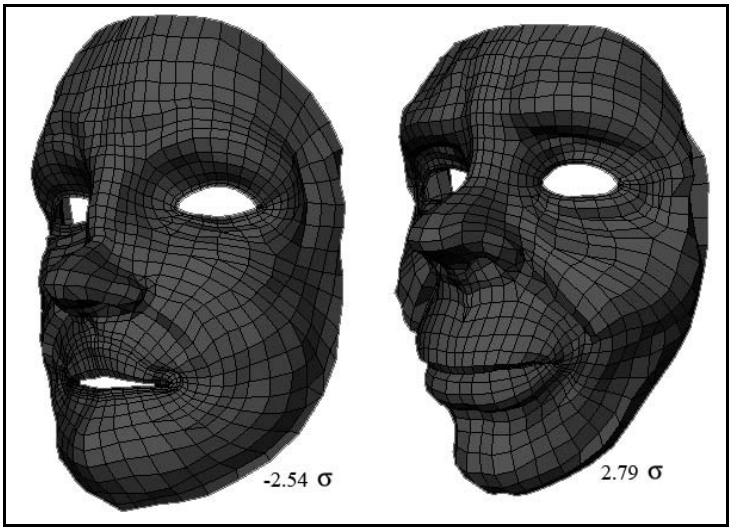

Figure 82: Varying the eleventh principal component coefficient of the neutral face set.

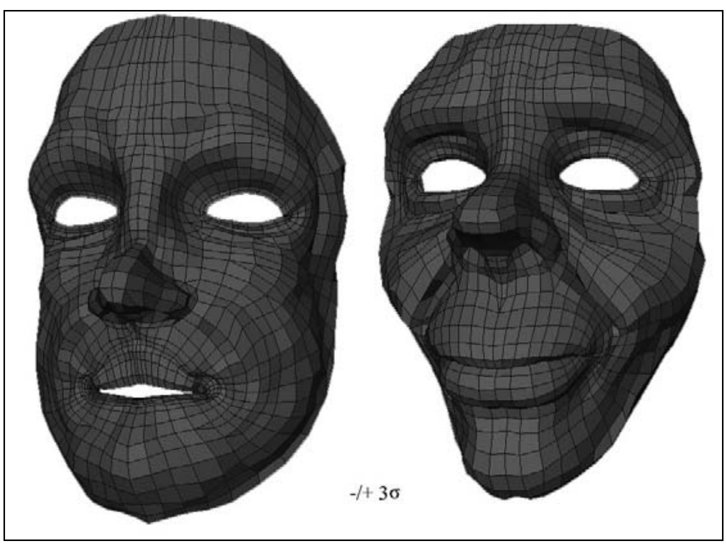

Figure 83: Varying the eleventh principal component coefficient of the neutral face set by -/+ 3 standard deviations.

Twelfth Principal Component

The twelfth principal component corresponds to the protrusion of the lips, nose, and cheek areas (Fig. 84-85). The majority of facial features (nose, cheek, and inner brow) are placed higher on the face.

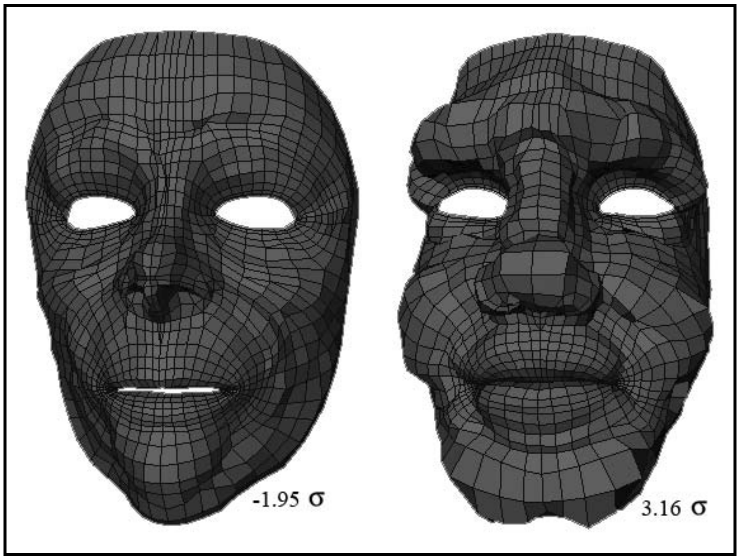

Figure 84: Varying the twelfth principal component coefficient of the neutral face set.

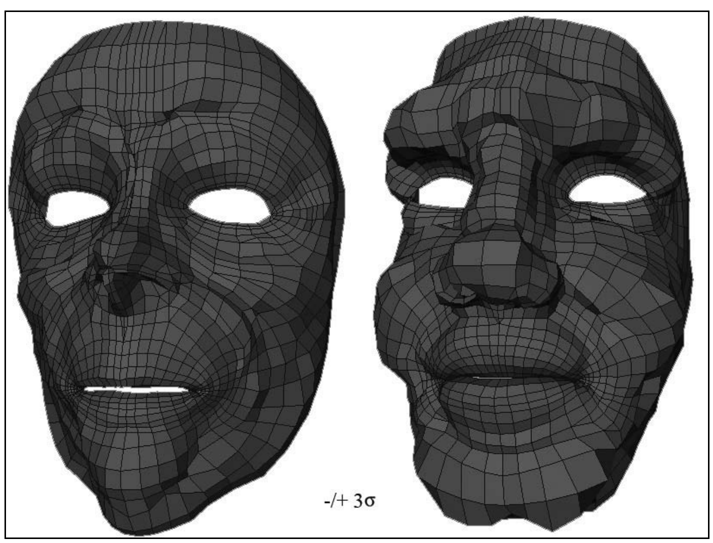

Figure 85: Varying the twelfth principal component coefficient of the neutral face set by -/+ 3 standard deviations.

Thirteenth Principal Component

The thirteenth component has an interesting relationship to the angle of facial features (Fig. 86-87). A negative coefficient resulted in an upturned brow, nose, lips, and chin. A positive coefficient angled all of these features downwards.

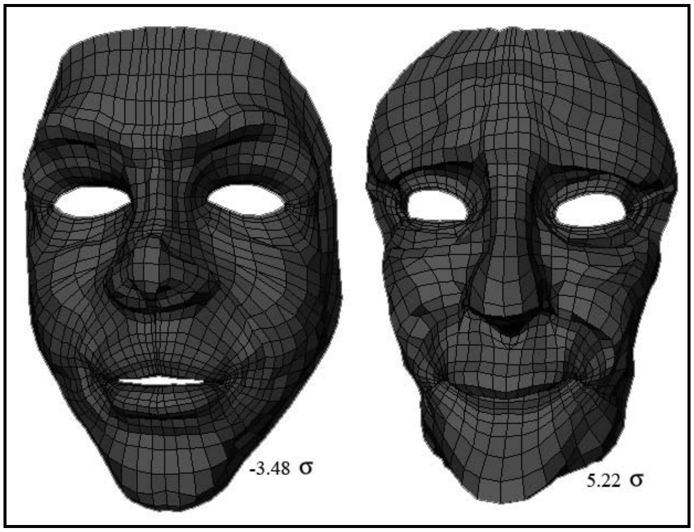

Figure 86: Varying the thirteenth principal component coefficient of the neutral face set.

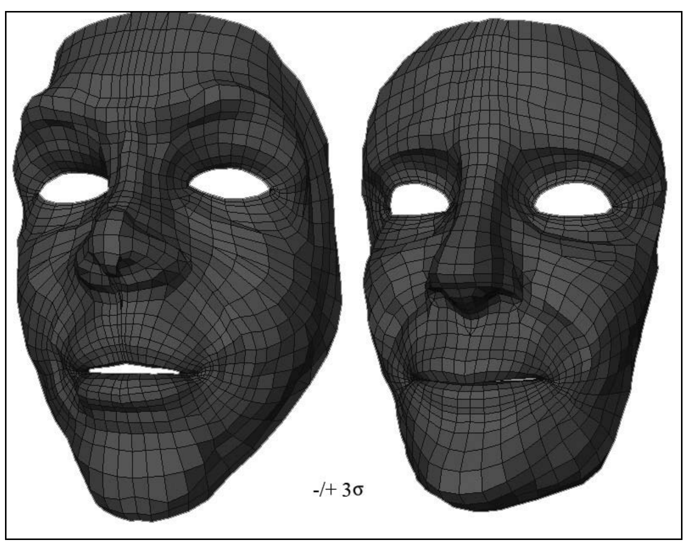

Figure 87: Varying the thirteenth principal component coefficient of the neutral face set by -/+ 3 standard deviations.

Fourteenth and Fifteenth Principal Components

Varying the fourteenth component of the neutral expression face set affects the cheek and brow areas (Fig. 88-89). A negative coefficient produces fatter, rounder cheeks and a high upturned brow while a positive coefficient creates gaunt, sunken cheekbones and a lowered brow. This change in brow angle affects the eye shape, making them wider and narrower. The fifteenth component has a similar correspondence to the cheek area, but also gives a variation in aspect ratio of the nose and chin (Fig. 90-91).

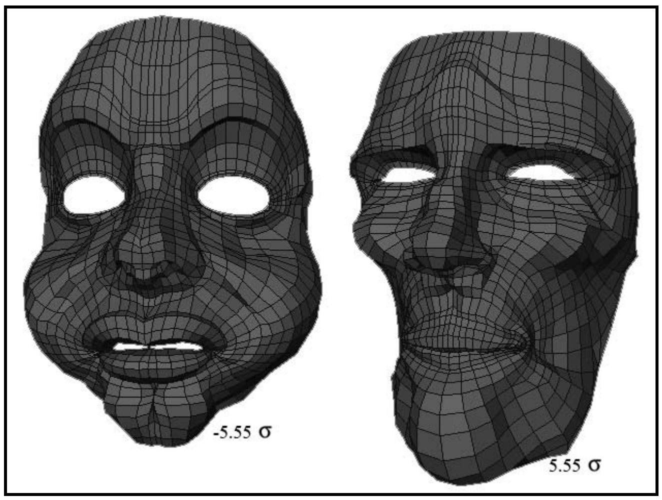

Figure 88: Varying the fourteenth principal component coefficient of the neutral face set.

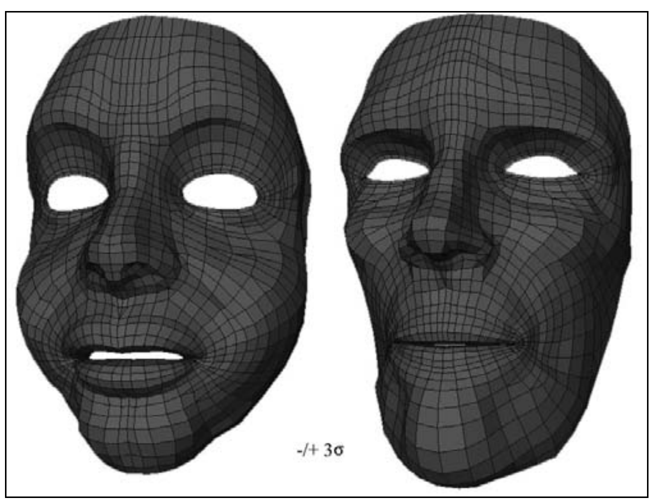

Figure 89: Varying the fourteenth principal component coefficient of the neutral face set by -/+ 3 standard deviations.

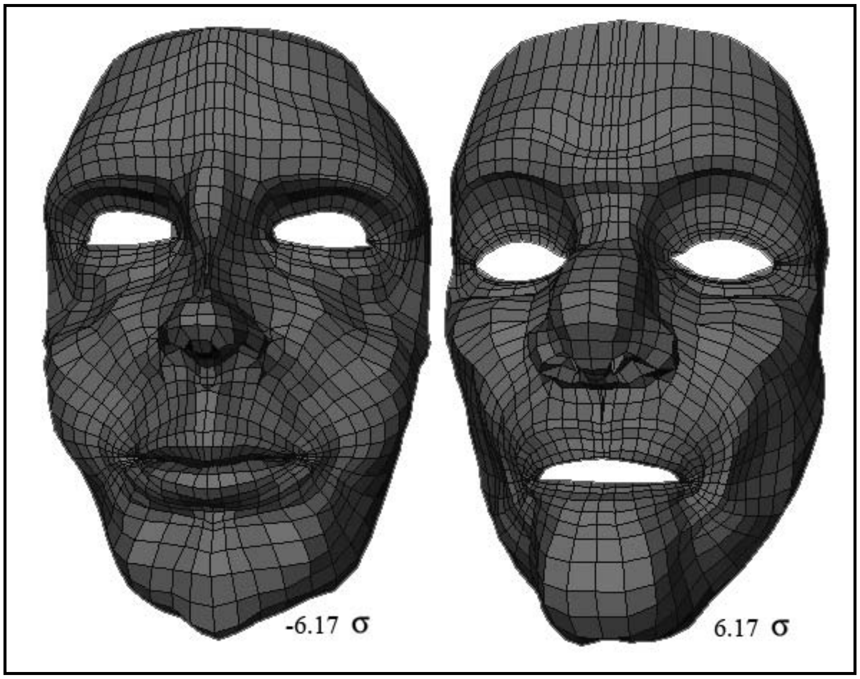

Figure 90: Varying the fifteenth principal component coefficient of the neutral face set.

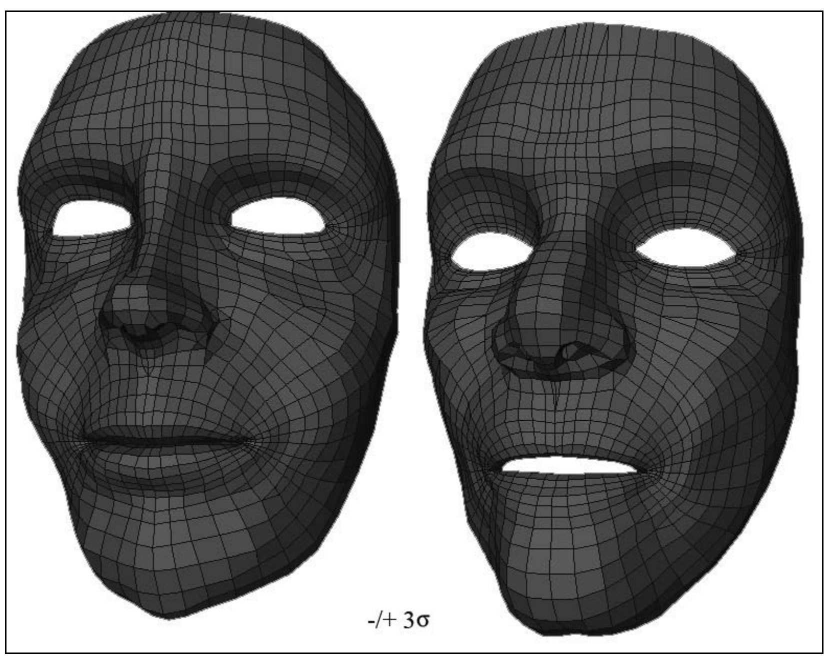

Figure 91: Varying the fifteenth principal component coefficient of the neutral face set by -/+ 3 standard deviations.

Observed Influence of Principal Components – Sorrow Face Space

Exploration of principal components of individual expression sets such as anger or sorrow did not separate expression from artistic style as in the neutral data. However the expression data sets guaranteed the creation of new faces with specific, non-neutral, expressions. Figure 92 shows the eigenvalues for the sorrow face space.

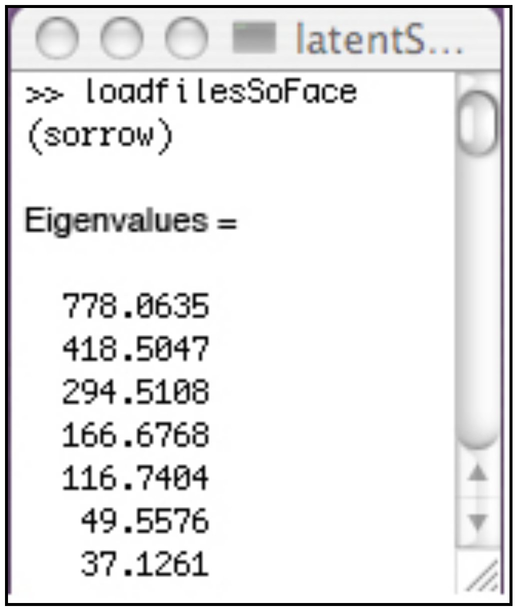

Figure 92: Screen capture of the eigenvalues for the sorrow face space.

First, Second, and Third Principal Components

The first, second, and third principal components of the sorrow expression face set act almost identically to the first, second, and third components of the neutral face data set (Fig. 93-98). The main difference is the presence of the sorrow expression and a smaller range of artist intended stylistic variation due to fewer artists represented.

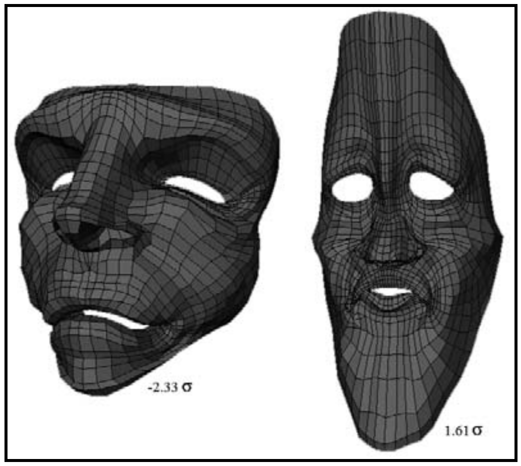

Figure 93: Varying the first principal component coefficient of the sorrow face set.

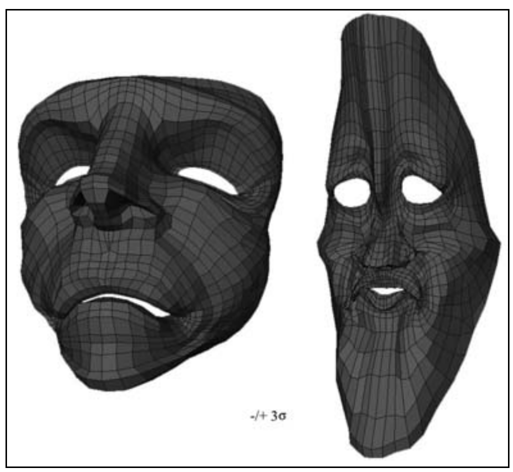

Figure 94: Varying the first principal component coefficient of the sorrow face set by -/+ 3 standard deviations.

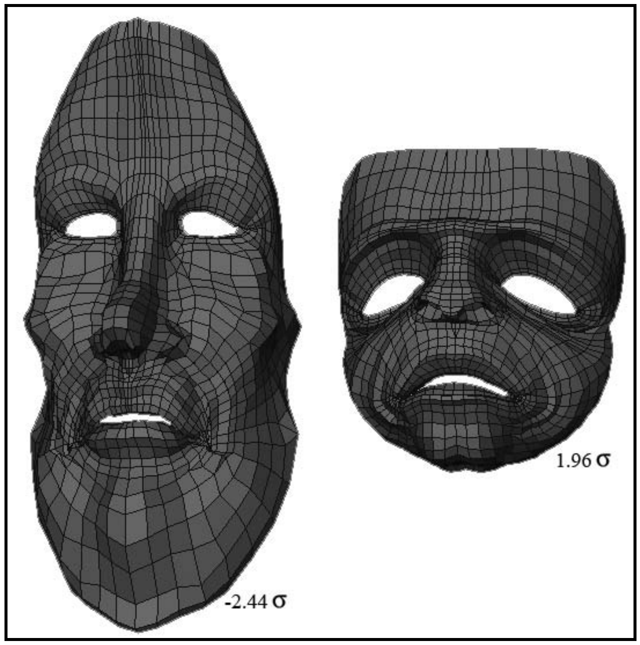

Figure 95: Varying the second principal component coefficient of the sorrow face set.

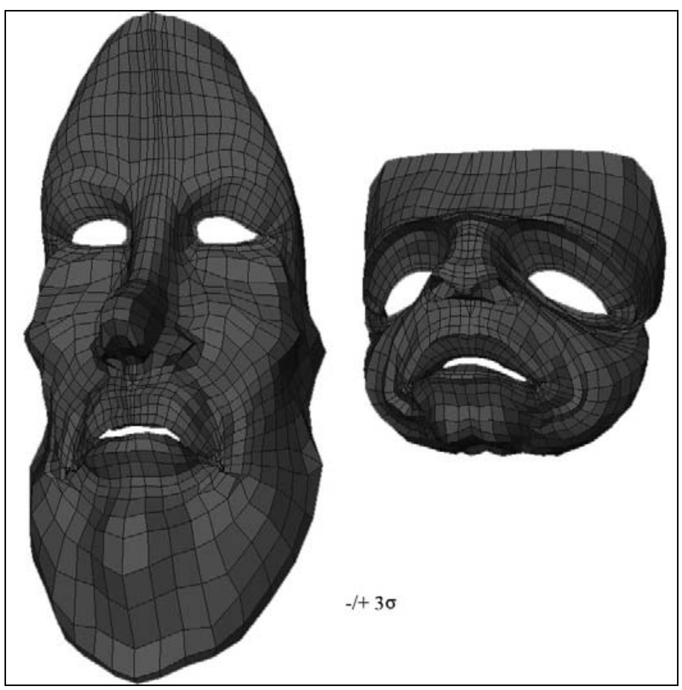

Figure 96: Varying the second principal component coefficient of the sorrow face set by -/+ 3 standard deviations.

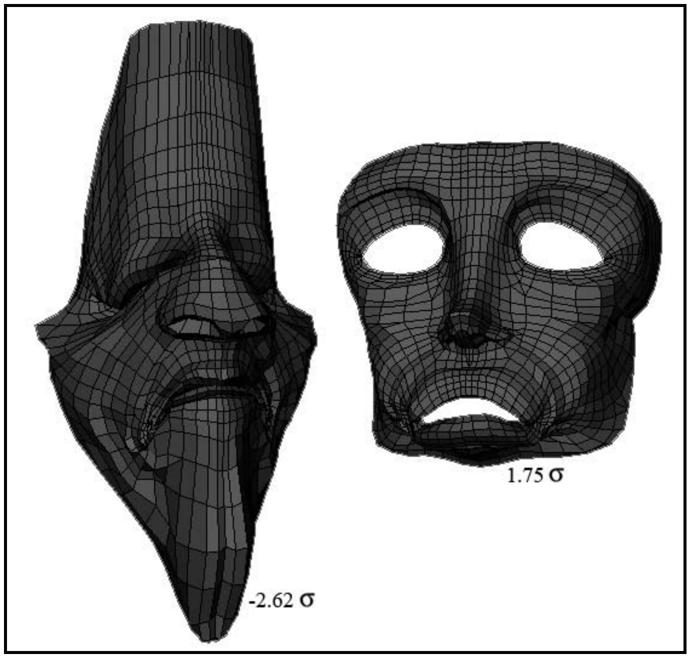

Figure 97: Varying the third principal component coefficient of the sorrow face set.

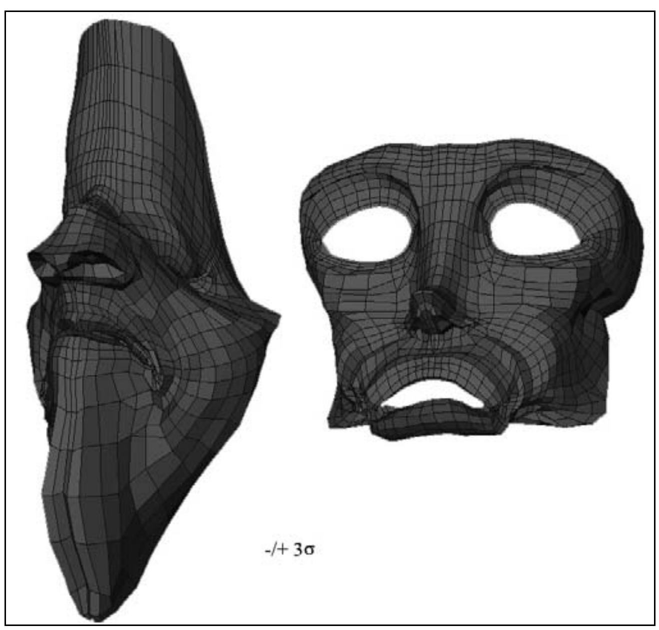

Figure 98: Varying the third principal component coefficient of the sorrow face set by -/+ 3 standard deviations.

Fourth and Fifth Principal Components

The fourth component corresponds to the protrusion of the brow, nose, lips, and chin. A change in turn of the nose was also identified (Fig. 99-100). Varying the fifth component resulted in a variation toward wide and open facial features, including the cheeks and chin (Fig. 101-102). A change in height of the cheek area was also identified.

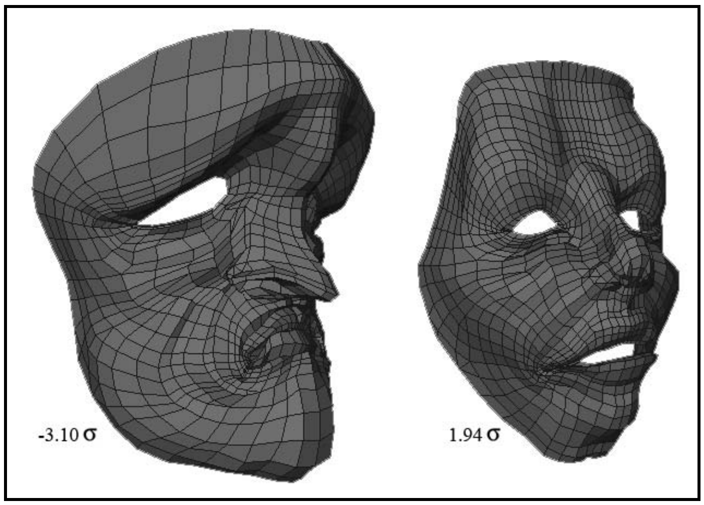

Figure 99: Varying the fourth principal component coefficient of the sorrow face set.

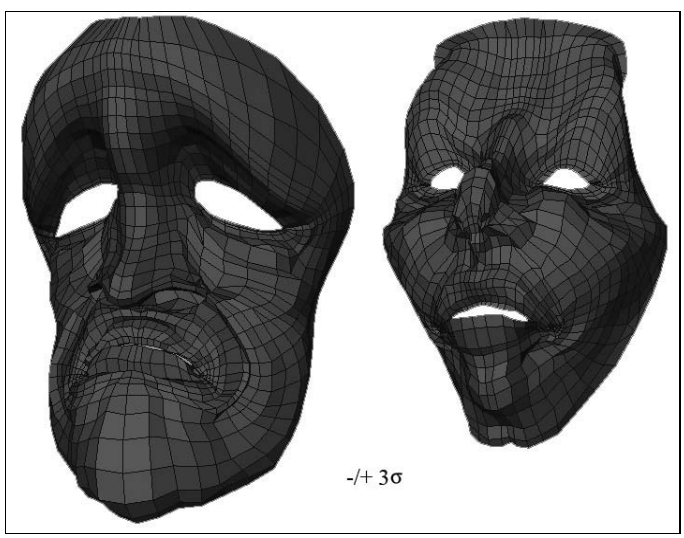

Figure 100: Varying the fourth principal component coefficient of the sorrow face set by -/+ 3 standard deviations.

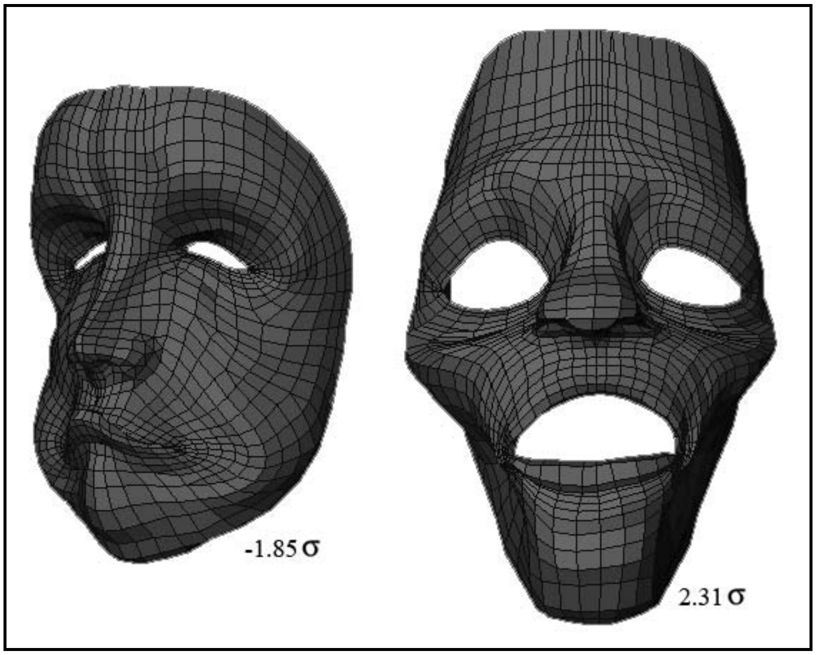

Figure 101: Varying the fifth principal component coefficient of the sorrow face set.

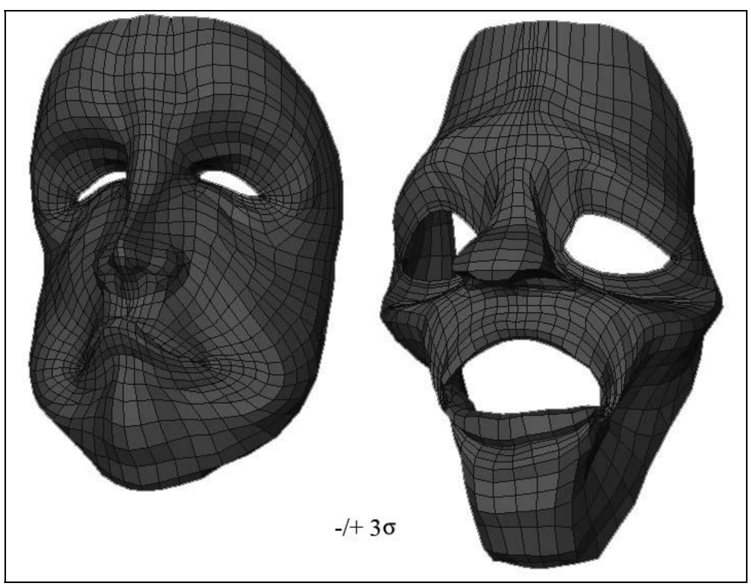

Figure 102: Varying the fifth principal component coefficient of the sorrow face set by -/+ 3 standard deviations.

Sixth Principal Component

The sixth principal component of the sorrow expression face set corresponded to changes to the eyes and corners of the mouth of the facial mesh (Fig. 103-104). Negatively scaling the sixth component created more upturned eyes and mouth corners. A positive coefficient resulted in more down-turned eyes and corners of the mouth. There was also a variation towards more protrusion of the facial features, including the ridge of the nose and area between the eyebrows.

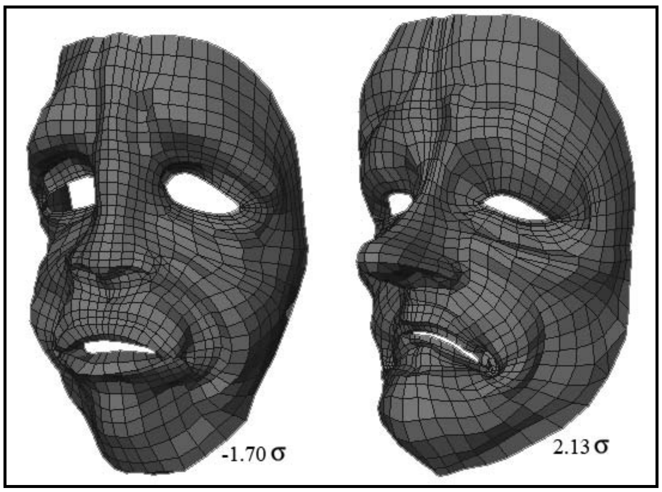

Figure 103: Varying the sixth principal component coefficient of the sorrow data set.

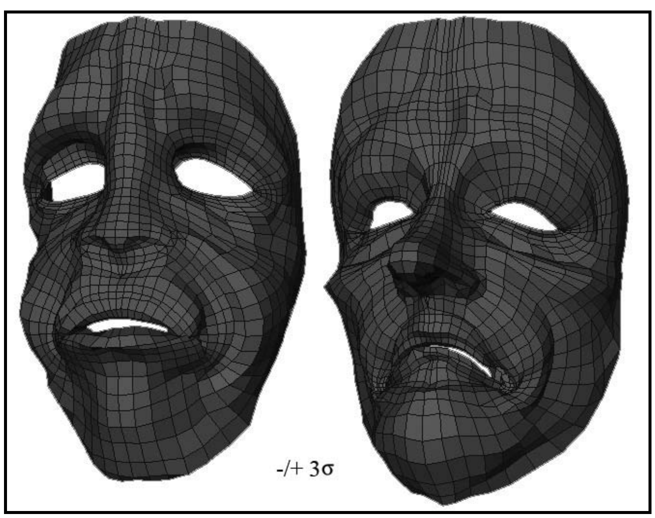

Figure 104: Varying the sixth principal component coefficient of the sorrow face set by -/+ 3 standard deviations.

Seventh Principal Component

The seventh, and final, component of the sorrow expression face set corresponded to the width of the nose and mouth, and the height of the inner brow (Fig.105-106). A negative coefficient created a wider mouth, narrower nose, and an upturned brow, while a positive coefficient created a narrower mouth, wider nose, and a lowered brow line.

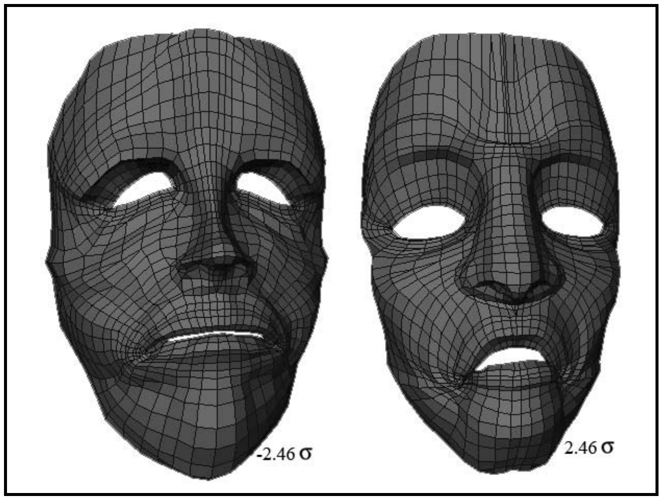

Figure 105: Varying the seventh principal component coefficient of the sorrow face set.

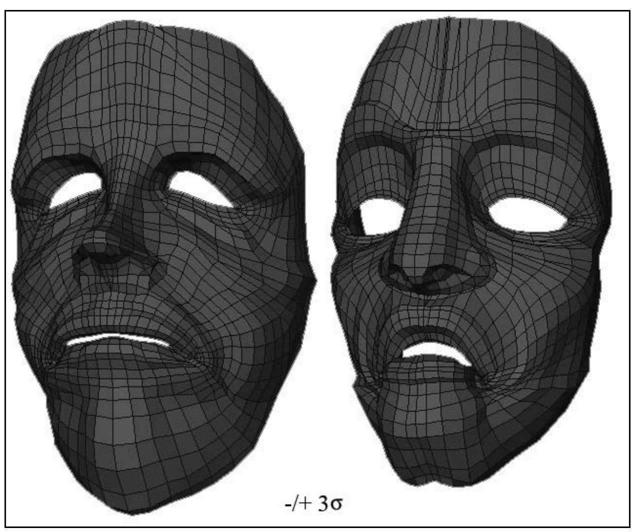

Figure 106: Varying the seventh principal component coefficient of the sorrow face set by -/+ 3 standard deviations.

Observed Influence of Principal Components – Other Face Spaces

Manipulating the principal components for the other expression sets results in new and different created faces. The types of properties and corresponding variation as seen in the previous data sets were also identified in these expression sets as well.

Anger Face Space

The complete set of eigenvalues for the anger created face space is shown in Figure 107. Figure 108 shows three created faces from the anger face space using combinations of these principal components. The left face was created by modifying the sixth and seventh principal component coefficients, resulting in a large chin, upturned nose, and protruding lips. The center face was created using only the second principal component and has broad nose, forehead, and chin areas. The face on the right was created using only the fourth principal component. It has a distinct triangular shape with a smaller nose and lips.

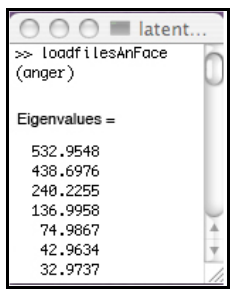

Figure 107: Screen capture of the eigenvalues for the anger face space.

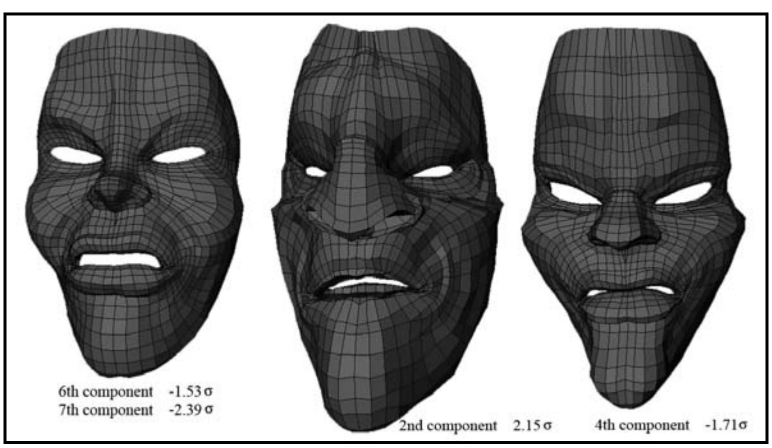

Figure 108: Faces created by varying principal component coefficients of the anger expression face set.

Joy Face Space

The complete set of eigenvalues for the joy created face space is shown in Figure 109. Figure 110 shows three created faces from the joy face space. The left face was created by modifying the seventh principal component coefficient, resulting in small squinting eyes and almost nonexistent lips. The center face was created by modifying the third and fourth principal component coefficients. There is a great protrusion of the lips and nose while the forehead and chin areas recede. The face on the right was created using only the second principal component. It has a broad forehead shape with a narrower chin and the nose is upturned.

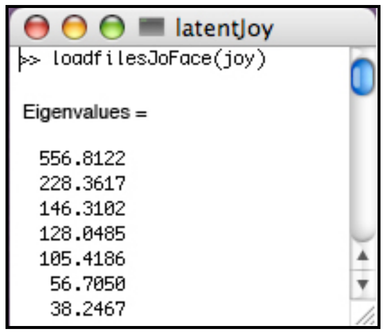

Figure 109: Screen capture of the eigenvalues for the joy face space.

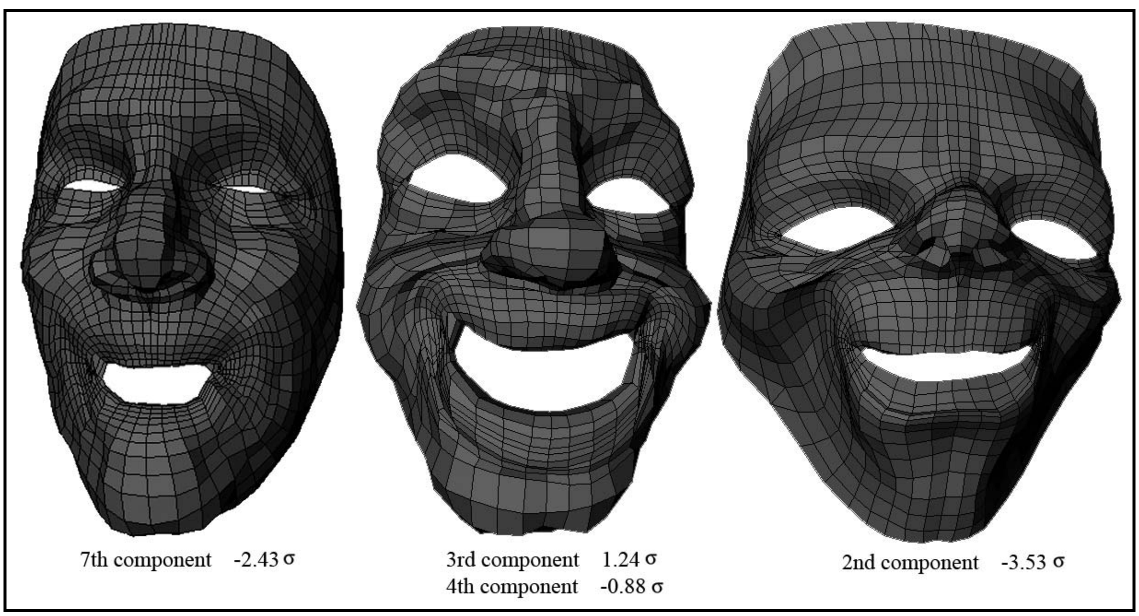

Figure 110: Faces created by varying principal component coefficients of the joy expression face set.

Disgust Face Space

The complete set of eigenvalues for the disgust created face space is shown in Figure 111. Figure 112 shows three faces created from different combinations of principal components. The left face was created by modifying the second principal component coefficient, resulting in large bulbous features for the forehead, lips, and nose. The center and right faces were created by modifying the sixth and seventh, and first and third principal component coefficients, respectively. While similar in overall shape, the features of the right face are higher and more skewed in expression than the center face.

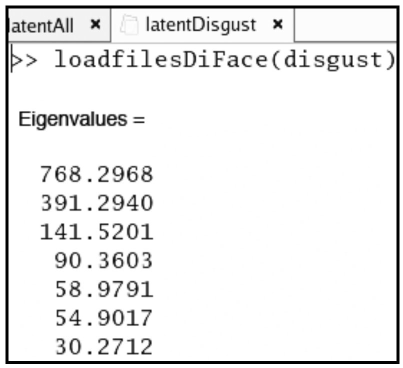

Figure 111: Screen capture of the eigenvalues for the disgust face space.

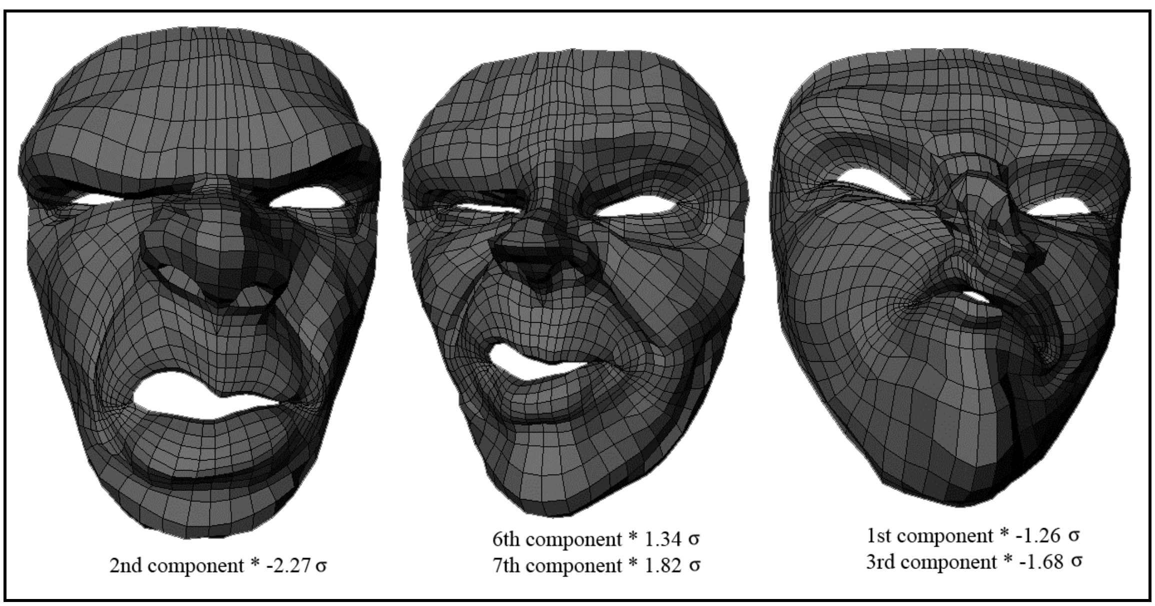

Figure 112: Faces created by varying principal component coefficients of the disgust expression face set.

Surprise Face Space

The complete set of eigenvalues for the surprise created face space is shown in Figure 113. Combinations of these principal components were used to created the three faces shown in Figure 114. The face on the left is created by modifying the first principal component coefficient. The lengthening of the face and widening of the mouth is similar to effects seen in other face spaces while using the first principal component. The center face is influenced by the eighth principal component. It has very high, round cheeks and thin lips. The right face was created by modifying the second and sixth principal component coefficients, resulting in the small nose and fuller lips.

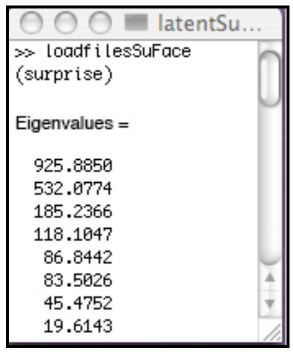

Figure 113: Screen capture of the eigenvalues for the surprise face set.

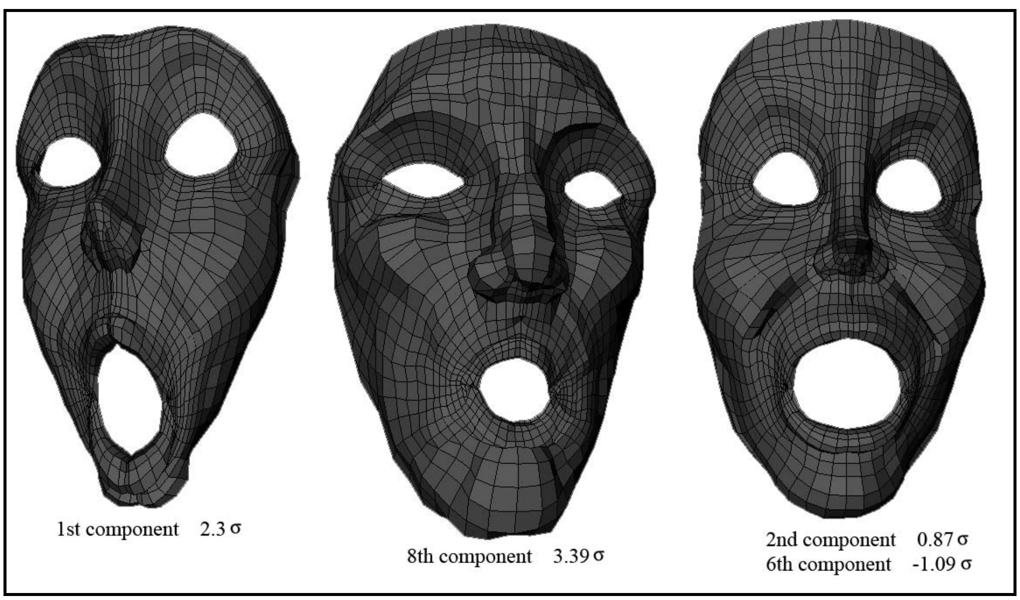

Figure 114: Faces created by varying principal component coefficients of the surprise expression face set.

Fear Face Space

The complete set of eigenvalues for the fear created face space is shown in Figure 115. (Note that the eigenvalues are multiples of 1000.) Figure 116 shows three created faces from the fear face space. The face on the left is influenced by the second principal component. This component varies the length and width of the face. The nose is high and down-turned. The center face is markedly feminine, with a pointy nose and full lips. It is influenced by the third principal component. Finally, the right face was created by modifying the fourth principal component coefficient. The dramatic turn of the corners of the mouth, high nose, and large nostrils are a result of this component.

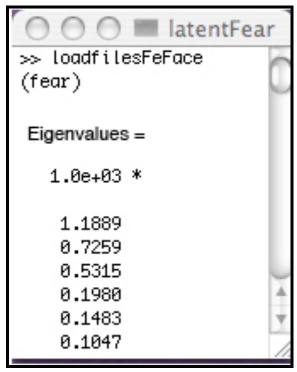

Figure 115: Screen capture of the eigenvalues for the fear face space.

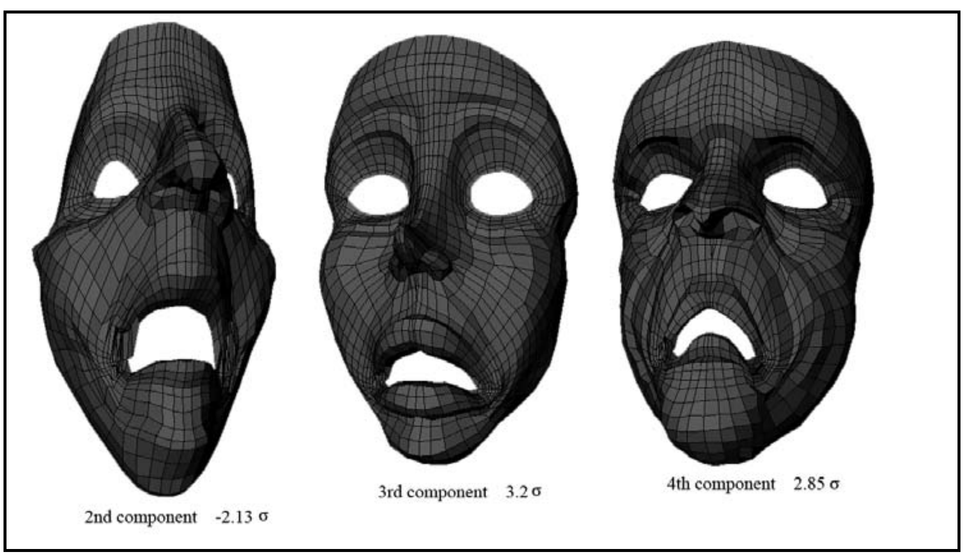

Figure 116: Faces created by varying principal components of the fear expression face set.

Overall Observations

Once these observations were made it was not difficult to create new faces with certain goals in mind. Aspect ratio, forehead width to chin width ratio, and expression could be controlled in a predictable manner. Some more general control of other features was possible. The size and placement of nose, lip, and cheek areas could be controlled, but altering them could have an unwanted affect elsewhere on the face. Once initial control choices are made, adding additional scaled eigenvectors provides unique and previously unknown choices for refining the appearance of the newly created face. I believe this work could be the basis for an interesting tool for artists.

Feature-based Transformations

The final objective was to create new faces based on specific characteristics such as gender, expression, and facial features. The goal was to use the previous principal component interface to create a new face and then add variation of specific characteristics. The nature of principal components does not allow control of these specific attributes in a predictable or individual manner, but using difference vectors representing variance of specific characteristics does.

I calculated several weighted averages of the entire face data set according to various attributes. Faces with the greatest presence of the attribute were scaled by 1, while faces without the attribute were scaled by 0. For a weighted average based on female features, a face which appeared female was given a coefficient of 1 while an ambiguous face was given a coefficient of 0.5. A male face would be weighed as 0 and therefore excluded. A weighted average subtracted from the average face for the entire data set formed a difference vector. This difference vector could be used to vary the amount of that attributes' influence on a created face. Difference vectors were calculated for each desired specific attribute. Scaling these vectors, in the same way the principal components were scaled, increased or decreased the amount of the attributes in the created face (Fig. 117 - 118).

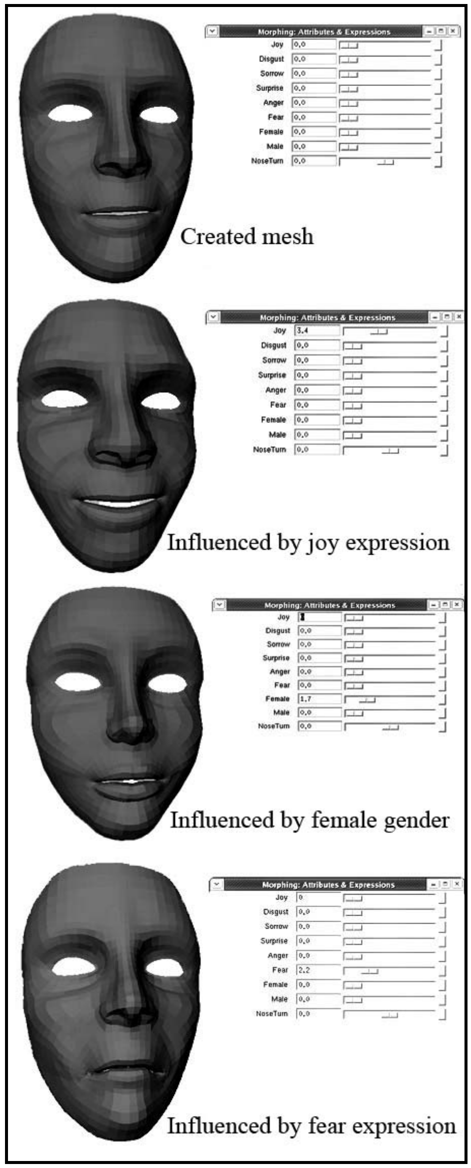

Figure 117: Created face influenced by various attributes.

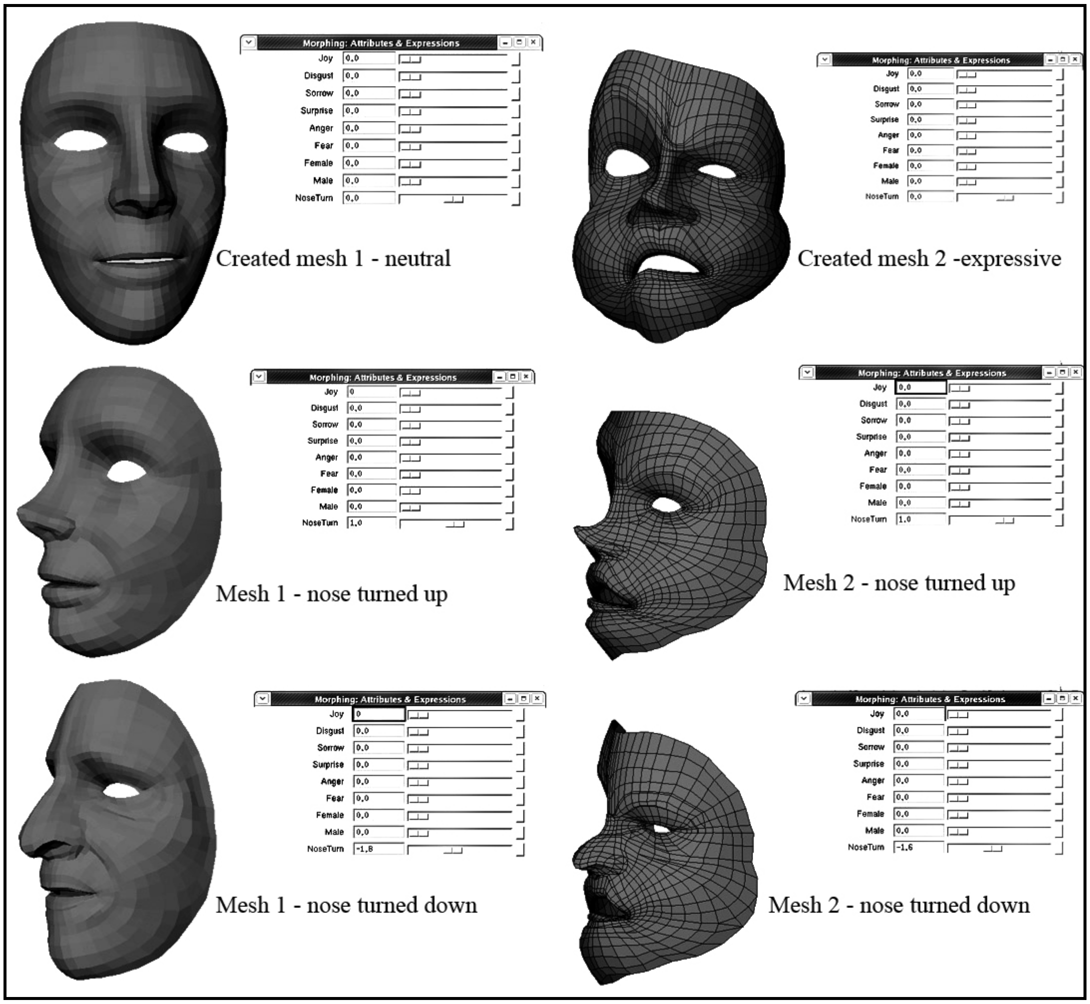

Figure 118: Created faces with modifications to nose turn direction.

This method worked, but the level of success depended on the original form of the created face being manipulated. A more neutral starting face could be successfully influenced by each attribute, while a created face that already had an obvious expression did not take on a new expression very well (Fig. 119). The results were also greatly influenced by the number of faces in the weighted data set. Controlling the femaleness of a created face was not as successful as controlling the joy of a face because there were fewer female faces in the weighted female average than there in the joyful average. A larger data set overall, with a better distribution of identifiable gender, and with every artist representing every expression, would likely increase the success of this method. This would be important if you needed more predictable results starting with any created face. This would also matter if you needed a greater range of variation of a specific characteristic without severely distorting the original created face.

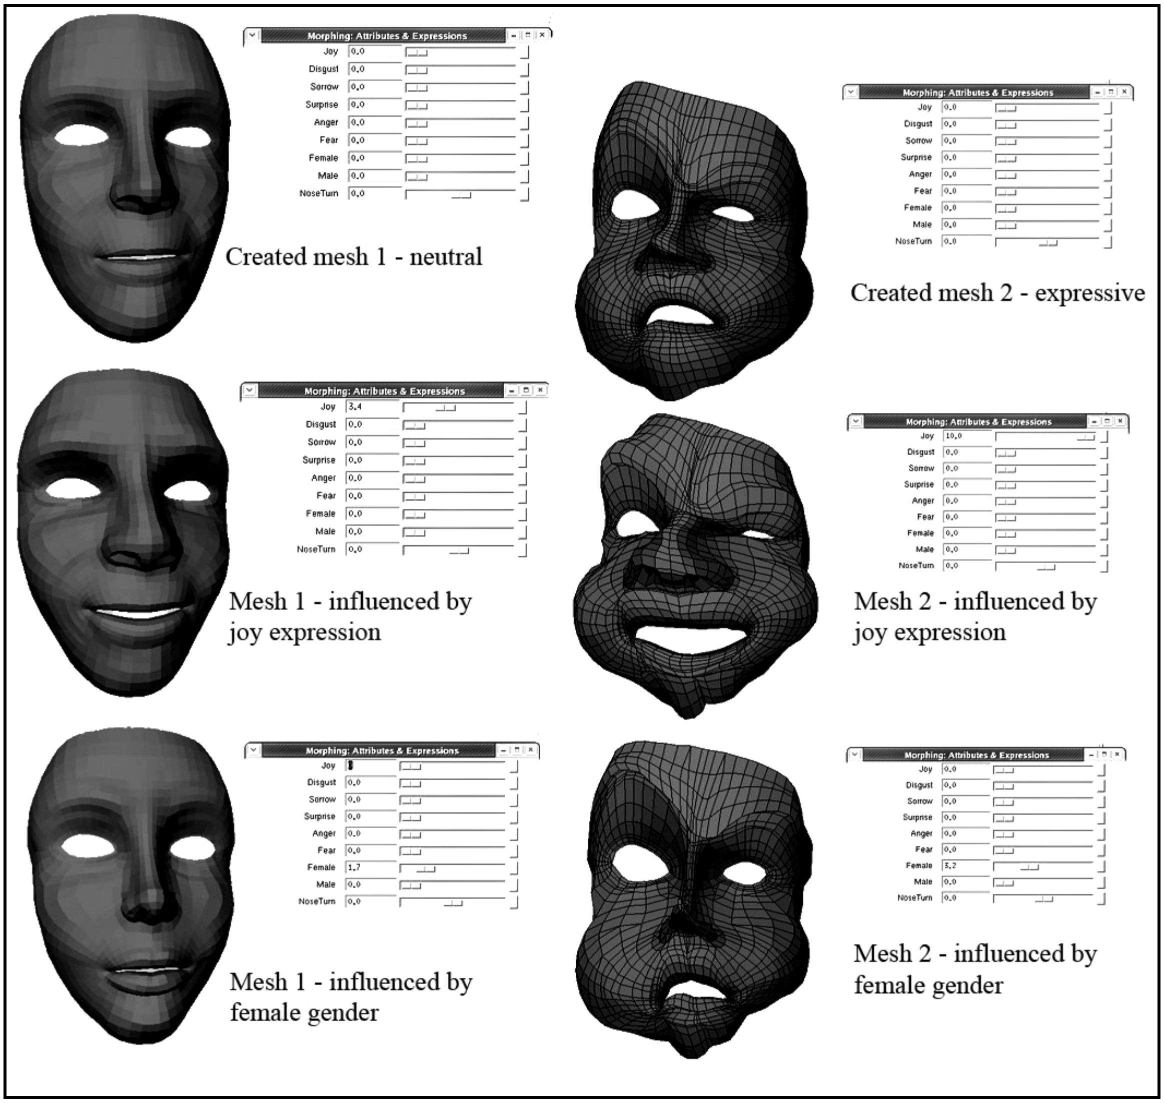

Figure 119: Comparing the influence of several attributes on two different created meshes.

CHAPTER VII

SUMMARY

The set of exemplar 3D face sculptures created by 16 volunteer artists enabled the development of a multi-dimensional vector space for describing faces. The faces included represent specific expressions (anger, fear, joy, surprise, disgust, and sorrow) and a sampling of artistic styles. The entire face set was used for analysis. It was also broken into separate face sets by expression. Principal components of the exemplar data were identified and the dimensionality of the vector spaces reduced by using only a limited number of components. This was done for the entire data set as well as its sub-sets.

New faces were created by modifying the average face of a data set by its principal components. These components were scaled by coefficient values and added to the average face of each data set.

The successful creation of a graphical user interface facilitated the interactive creation and viewing of these faces. The user could then interactively adjust the values of the principal component coefficients used. This interface also provided the ability to create animations of these faces morphing between previously specified faces. Viewing those animations enabled a better understanding of how the influence of various principal components changed a created face.

Feature-based controls allowed the variation of specific expression, facial features, and characteristics such as gender. Those controls allowed the further transformation of a face created using the principal component interface. These feature-based transformations were based on the use of difference vectors, not principal components. A second graphical user interface was developed for these feature-based controls.

Conclusions

I had thought that I would be able to predict the effects of varying certain components to obtain certain characteristics. However, principal components alone could not easily control independent features and characteristics due to the generalized nature of the analysis. Although I could often predict how one feature would be affected by a specific component coefficient variation, other features would be affected as well. Combining variations of multiple component coefficients further reduced the predictability of the results. For example, specifying a non-zero value for the first principal component coefficient as well as a non-zero value for the second principal component coefficient would have a specific result. However, this pairing would dramatically change the results of further variation the first principal component coefficient. Though unpredictable, the exponentially large number of possible combinations of grouping multiple principal components presented interesting design choices for creating new faces. Each design choice would then determine future artistic choices for each created face. The design choices available would probably be different if a different group of artists, and therefore a different set of exemplar faces, were used. While the first few components might represent similar data variance, such as aspect ratio and size of the face, less significant components would likely contain different and unique data variance.

Any face created could then be used as the base facial mesh for the feature-based transformation interface. It then was possible to overlay specific features and characteristics onto faces previously created using the principal components. While fairly successful, the results from the feature-based method might be improved by using a larger exemplar data set with a greater range of desired characteristics such as gender or facial feature. Because this method used difference vectors as a basis for transformation, the results seem more closely linked to the shapes within the exemplar data set. Overall, the created graphical user interfaces, along with the data extracted from the artist created exemplar faces, seem to be the basis for a useful tool for artists.

CHAPTER VIII

FUTURE WORK

One extension of this work would be to combine principal component analysis of the data set with individual component analysis for more direct control over facial features when creating new faces. Independent component analysis (ICA) is a statistical computational technique for revealing factors that underlie sets of random variables or measurements by linearly extracting independent features from the data [HyvarinenHoyerIniki00]. This might allow separation of expressive information from artistic style. A combination of PCA and ICA might be a powerful tool in understanding the control of, and uses for, a virtual face space.

This work explored aspects of virtual faces created by a certain number of artists . Another option would be to focus on a specific artist rather than many, using either method of statistical analysis. Using only various works from one artist might allow identification of the salient aspects of his or her unique style. Such a study could use created virtual face sculptures as did this work, or could use digitized physical sculpture or pre-existing digital models.

Another possibility would be to apply these techniques to two-dimensional artistic works. New criteria for data gathering and measurement would have to be defined based on the characteristics of such works. Possible works might include digital or analog drawings, paintings, and photographs. This data could be from many artists for a more general study, or only one artist for more specific exploration of his or her style.

These future explorations could focus on representations of faces or on more general themes. Explorations could be extended to non-human or humanoid faces such as those belonging to animals or creatures of mythology, religion, or fiction. Future work could focus on other recognizable aspects of a character such as hands, feet, or entire bodies. In addition, non-character related studies might focus on inanimate objects or abstract sculptures.

Finally, my work might provide the basis for developing a useful tool for artists. Such a tool could include a more extensive interface, texturing abilities, or advanced animation controls for the created faces.

REFERENCES

BLANZ, V. AND VETTER, T. 2000. A morphable model for the synthesis of 3D faces, In *Proceedings of ACM SIGGRAPH 2000*, ACM Press / ACM SIGGRAPH, Computer Graphics Proceedings, Annual Conference Series, ACM, 187-194.

EKMAN, P. 1989. The argument and evidence about universals in facial expressions of emotion, *Handbook of Social Psychophysiology*, H. Wagner and A. Monstead, Eds., John Wiley, Chichestor, Great Britain.

HYVARINEN, A., HOYER, P. O., and INIKI, M. 2000. The Independence Assumption: Analyzing the Independence of the Components by Topography. *Advances in Independent Component Analysis*, M. GIROLAMI, Eds., Springer-Verlag, London, Great Britain.

HILL, T., AND LEWICKI, P. 2006. *Statistics: Methods and Applications, a Comprehensive Reference for Science, Industry, and Data Mining*, StatSoft Inc., Tulsa, Oklahoma.

PARKE, F. I. 1972. *Computer Generated Animation of Faces*, M.S. Thesis, University of Utah, Salt Lake City, UT.

PARKE, F. I. 1974. *A Parametric Model for Human Faces*, PhD Thesis, University of Utah, Salt Lake City, UT.

PARKE, F. I. 1991. Techniques for facial animation. *New Trends in Animation and Visualization*, N. Magnenat-Thalmann and D. Thalmann, Eds., John Wiley, Chichester, Great Britain.

PARKE, F. I., AND WATERS, K. 1996. *Computer Facial Animation*. A. K. Peters, Ltd.

PLATT, S. M., AND BADLER, N. I. 1981. Animating Facial Expressions, In *Proceedings of ACM SIGGRAPH 1981*, ACM Press / ACM SIGGRAPH, Computer Graphics Proceedings, Annual Conference Series, ACM, 245-252

WATERS, K. 1987. A muscle model for animating three-dimensional facial expressions. In *Proceedings of ACM SIGGRAPH 1987*, ACM Press / ACM SIGGRAPH, Computer Graphics Proceedings, Annual Conference Series, ACM, 17-24

VITA

Name: Jessica Lauren Riewe

Address: Texas A&M University C418 Langford Center 3137 TAMU College Station, Tx 77843-3137

Email Address: jriewe@viz.tamu.edu

Education: B.S., Environmental Design, Texas A&M University, 2004
M.S., Visualization Sciences, Texas A&M University, 2007