Model-Based Face Computation

1. Research Team

Project Leader: Prof. Ulrich Neumann, *IMSC and Computer Science*

Post Doc(s): John P. Lewis

Graduate Students: Hea-juen Hwang, Zhenyao Mo, Gordon Thomas

2. Statement of Project Goals

Prior knowledge of the canonical structure of the human face can aid in various automated face-processing tasks. In this project we have developed a statistical appearance model for faces and are exploring its application to several problems: stylized face rendering, caricature, and reconstruction of occluded face images.

3. Project Role in Support of IMSC Strategic Plan

Model-Based Face computation is part of the general IMSC effort towards expressive human interaction in virtual and augmented reality environments. While this project involves processing based on the *structure* of the face, the complementary Data-Driven Facial Animation project is directed toward deriving models of facial *movement* (including non-speech gestures) directly from data.

4. Discussion of Methodology Used

A two-dimensional statistical prior model of faces was constructed using the principal components of the shape and texture of a large number of face photographs. Specifically, 94 feature points were manually located on 500 faces selected from the FERET database [1]. The shape variation is then removed by morphing all faces to the mean shape and a second model describing the face texture variation is obtained.

This shape+texture linear subspace modeling approach was independently pioneered by a number of research groups and goes by various names including Active Appearance Models and Morphable Models. Our contributions include:

- *Group modeling and outlier rejection.* A comprehensive facial model should be able to represent faces of both sexes and of different ages and ethnicities. On the other hand, if abundant facial data for young Asian females (for example) is available, adding additional data on (for example) older black males will not help much in improving the data model for the young Asian female, and indeed it may degrade the model. In our work we replace the global principal component analysis with a representation that considers near neighbors to be part of the same "group". This enhances modeling and prediction even given limited data (see Figure 1).

- *Model extrapolation*. We use the statistics implicit in the data to extrapolate new face regions. In Figure 1, regions of novel faces not among the training pictures (left) are removed (center). The model then predicts the missing regions (right). The prediction is not always accurate, but if there are any statistical regularities in the data they will be revealed by our process. For example, if a particular "type" of face almost always goes along with a wide mouth and thin lips, then the prediction will be of a wide mouth and thin lips. If the statistical trends are weak then the extrapolated result will appear blurry, as in the second row of Figure 1.

Figure 1. Regions of novel faces not among the training data (left) are obscured (center). Our model then predicts the missing regions (right).

5. Short Description of Achievements in Previous Years

This is a new project initiated in the past year, although the statistical face model was developed as part of the Data-Driven Face Modeling project of 2003.

5a. Detail of Accomplishments During the Past Year

Two examples of model-based face computation have been developed:

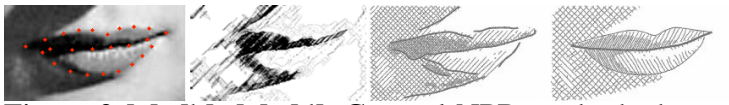

Figure 2 [a], [b], [c], [d]. General NPR methods do not produce good line portraits of faces: [b] a Photoshop filter creates strokes that cross the line between the lips; [c] strokes derived from edge detection cannot distinguish shadow edges from facial features; [d] this drawing is rendered consulting the model, that is, the mouth shape (depicted in [a] as red points).

Automatic Line-Drawing Portraits. Existing non-photorealistic rendering (NPR) techniques produce poor portraits in some cases. This is particularly true when the portrait is simple, containing only a few lines. General NPR techniques cannot produce this sort of portrait because they do not "know" which lines are important and which can be discarded (Figure 2). By comparing the luminance to a prior face model (fit to the particular face) we are able to assign salience to various possible lines, and keep only the most important ones (Figure 3). Several of the resulting automatic portraits are shown in Figure 4.

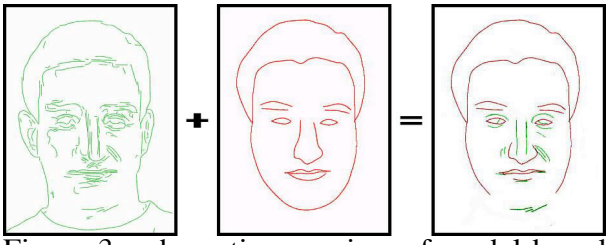

Figure 3. schematic overview of model-based portrait generation: an image-based description of a face (left) is reduced to face features (right) by consulting the model (middle).

Caricature. S. Brennan formalized caricature by comparing the proportions of a particular face to those of a mean face, and exaggerating any differences found [2]. In this sub-project we improve Brennan's approach in several respects:

- The population variance is taken into account. Features that are different from the mean but are well within the population variance are not exaggerated.
- We work with photographic textures rather than simple line drawings.

Figure 5 shows some initial results of automatic caricature.

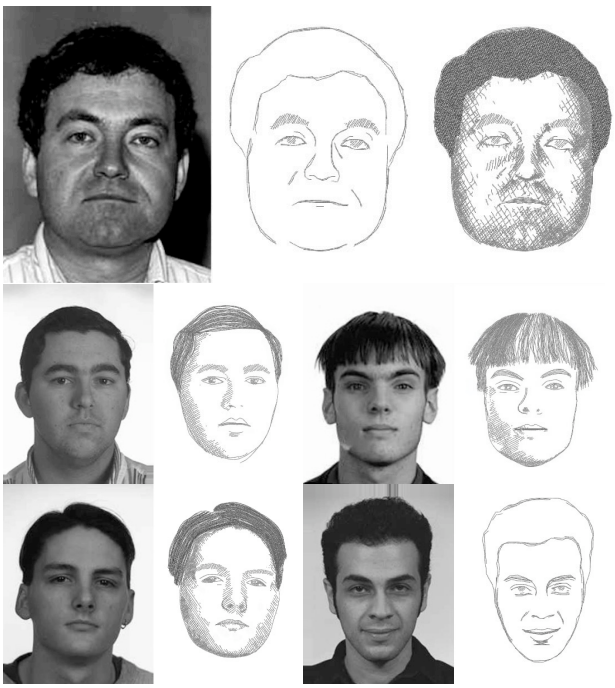

Figure 4. Some portraits automatically produced using our method. The two portraits at top show how the portrait complexity can be varied with the model-based approach.

6. Other Relevant Work Being Conducted and How this Project is Different

Related problems have been explored at the Microsoft Research China by H. Shum and collaborators. Their approach considers the nature of caricature and portraiture to be artist-specific, and as such they use machine learning to mimic the work of a particular artist. While this produces good results, it does not easily generalize beyond the orientation and style chosen by the particular artist. Our work explores the idea that both portraiture and caricature can be described with a general model-based theory. This in turn allows different styles to be obtained programmatically and without additional data gathering (Figure 4).

Figure 5. Caricatures automatically produced using our method

7. Plan for the Next Year

The recent results clearly show the benefit of introducing a detailed prior face model. This model is currently two-dimensional, however, which limits its application to near-frontal poses. We are beginning an effort to develop a corresponding three-dimensional face model.

8. Expected Milestones and Deliverables

The current deliverables include several hand-annotated facial databases, and the two prototype programs described above (roughly 10,000 lines of Java). Our future milestones are indicated in the following table:

- 2D statistical face prior model
- Demonstration of model-based portraits and caricature
- Develop 3D face modeling system
- If successful, develop 3D face database and prior model

9. Member Company Benefits

Sophisticated statistical models of face appearance such as the one we prototyped have a number of uses, including:

- Older or younger versions of a face can be approximately extrapolated, by first fitting the face using the model and then moving along an age axis in “face space”.
- The estimation of a missing face part may have application in reconstructive surgery: in cases where a face region is deformed at birth, there is no prior photograph to provide a goal for the surgery. As in the previous item it should be possible to define a “beauty axis” in face space given human ratings of the beauty of diverse faces. With this

definition, it would also be possible to provide several reconstructive surgery candidates rather than one "mean" choice. The several choices would be situated at several points along the beauty dimension, while still being consistent with the surrounding face data as in Figure 1.

- Similarly, the estimation of missing face parts may be useful in security or other person identification contexts, as when a photograph of an individual is available but portions of their face are obscured.

Existing face-finding algorithms often cannot localize parts of a found face other than the eyes and mouth center. By pairing such an algorithm with a statistical appearance model it will be possible to automatically locate detailed facial features such as the mouth corners, nostrils, points along the eyebrows, etc. This capability is in turn useful in diverse areas including previewing medical procedures, virtual cosmetic preview, and various computer gaming applications.

10. References

[1] P. J. Phillips, H. Wechsler, J. Huang, and P. Rauss. The FERET database and evaluation procedure for face recognition algorithms. Image and Vision Computing J., 16(5):295–306, 1998.

[2] S. Brennan, Caricature generator, Master's thesis, MIT, 1982.