Animating Facial Expressions

Stephen M. Platt
Norman I. Badler

Department of Computer
and Information Science
University of Pennsylvania
Philadelphia, PA

Abstract: Recognition and simulation of actions performable on rigidly-jointed actors such as human bodies have been the subject of our research for some time. One part of an ongoing effort towards a total human movement simulator is to develop a system to perform the actions of American Sign Language (ASL). However, one of the "channels" of ASL communication, the face, presents problems which are not well handled by a rigid model.

An integrated system for an internal representation and simulation of the face is presented, along with a proposed image analysis model. Results from an implementation of the internal model and simulation modules are presented, as well as comments on the future of computer controlled recognition of facial actions.

We conclude with a discussion on extensions of the system, covering relations between flexible masses and rigid (jointed) ones. Applications of this theory into constrained actions, such as across rigid nonmoving sheets of bone (forehead, eyes) are also discussed.

Introduction: Representation and simulation of gross motions of the human body have been investigated [1],[2],[3], modelling the body as rigid segments, and controlling it with simulated processors placed at each joint. This concept has been extended in a human movement recognizer described by O'Rourke [11], which uses constraints placed on the figure by the joint actions of the interconnecting rigid limbs. These techniques, however, cannot be easily extended to modelling non-rigid masses. The work presented here deals with the problems involved in the manipulation of such deformable objects. The solution proposed incorporates the actual causes (motivators) of the actions, rather than just simulation the resulting actions directly.

The basic goal of this research is to devise an efficient and accurate model of the human face, and to develop or adapt a notational system to encode actions performable on the face. This notation system should lend itself not only to ease of reproduction, but should also be usable as a human entry system, and be sufficiently rich and well-defined to make computer recognition (via computer controlled camera) possible.

This research forms a part of a larger project involving the recognition and simulation of American Sign Language (Ameslan or ASL). As described in Baker [4] and Baker and Padden [5], ASL communication can be broken down into five channels: (1) the hands and arms, (2) the head, (3) the face, (4) the eyes, and (5) the total body posture. Our approaches to simulating the above items, with the exception of (4), are described in Badler et. al. [1].

We have discovered [1], [3] that it is usually more convenient to maintain two separate notations for body and action representation, and maintain a translation scheme between them. Action-based notations are the lower level notation -- they are more oriented towards the physical image of the object, as in a camera image. Structure-based notations are of a higher level -- they are based around the logical structure of the object, and are more suited to object manipulations.

As such, the facial action recognition to simulation system was divided into three logical sections, a camera processor, an internal model manipulator, and a face simulator.

This research was supported in part by NSF grant MCS78-07466.

```
        Input Image
             |
             |
             v
      Camera Processor
             |
[High Level Representation]
             |
             v
       Internal Model
        Manipulator
             |
 [Low Level Representation]
             |
             v
         Simulator
             |
             v
        Output Image
```

Fig. 1: The Basic System (general outline).

Initially, all that is available is an input from a computer controlled camera. The Camera Processor would manipulate this data, determining which facial actions are being performed and interpreting them in terms of the motion representation notation. This output (the motion representation notation) is passed to the Internal Model Manipulator. This module takes the motion representation and prepares it for simulation by converting it into the structures needed by the simulator. Finally, the last module, the Simulator, uses this data to manipulate its model of the face, producing suitable output.

Partitioning the processing in this fashion results in two tie-points where the data has been partially processed into a compact form -- both before and after the Internal Model Manipulator. Processing may be halted at either point, allowing storage of the image for future use, additions/modifications, or transmissions to distant sites for reconstruction.

Action-based Notations: We examined several action-based notation systems for applicability to the problem. Important features for such a system include ease of (computer) adaptation, completeness of the notation, and the system's extensibility to more general problems.

Labanotation [10] was examined first as it has already been used in several human motion studies [3], [2], and [7]. Labanotation is basically a dance notation scheme, used to record basic human motions. However, there is little in the notation for the recording of facial actions.

Another notation system is Sutton Notation [14]. Sutton Notation is more pictorially based, but still records actions as changes in body position on a staff. However, unlike Labanotation, Sutton Notation also has the capability to notate facial expressions. Sutton also mentions a project concerning the computer implementation of this notation.

Birdwhistell [6] describes another scheme for a communication-notation system. This system was designed to record the detailed motions and actions of body language and movement in interpersonal communication. It has a large vocabulary of (pictorial) symbols for many actions and positions of the body and face, but lacks a scheme for generalization of interactions between facial gestures.

The last system considered was the Facial Action Coding System (FACS) [8]. Unlike the previously described systems, FACS is not graphically oriented. Rather, it describes the set of all possible basic actions performable on a human face. Some sample actions are Inner Brow Raiser, Outer Brow Raiser, and Lid Tightener. Each basic action (called an Action Unit, or AU) is a minimal action in the sense that it cannot be broken up into smaller actions. AUs interact in several different fashions to build up a full facial expression.

The AUs are designed to be closely connected with the anatomy of the face. Each AU is caused by a minimal number of muscle contractions or relaxations. This further strengthens its claim to minimal actions since each AU is controlled by either a single muscle or a small set of closely related muscles.

The description of a facial expression using FACS is vastly different from descriptions in any of the previously considered systems. The notations of those systems are symbolic or pictorial, and while this aids the human user by providing a highly mnemonic symbology, it tends to make computer storage of the notation difficult. FACS, on the other hand, describes for each frame (facial expression to be analyzed) a list of the names of AUs being performed. This format is much more suitable to computer usage, although it increases the learning difficulty for human notators.

FACS itself is not easily extensible to other domains -- the definitions of the AUs are tied intimately to the human face. However, the theory behind FACS, notating based on minimal performance units (in turn based on anatomical structure), can be applied to other domains.

Structure-based Representations: As previously mentioned, there is a second level of internal representation: that of the actual face. It is at this second level that the action representations are interpreted and the simulation is performed. The internal structure of the face's representation must be well chosen, since the quality of any images produced will rest on this representation. For ease and accuracy of translation, the representation must also be easily accessed from the facial action coding.

The three techniques considered below can all be thought of as variations of a single representation -- simulating the face by a patchwork model. They only vary in the complexity of the techniques used to perform the facial actions.

The first of the models is the simple 2D surface patch technique. In this model, the head is broken into small patches of 'skin', as suggested by [9]. A facial action under this system would consist of simple warping of a subset of the set of skin patches:

```
Facial Action ==
     ((Skin-Patch-Number Trans Rot)
      (Skin-Patch-Number Trans Rot)
               .
               .
               .
```

where Skin-Patch-Number is an individual patch identifier, Trans is the triple (dx dy dz) -- a 3D translation, and Rot is the triple (rho theta phi) -- the rotation the patch undergoes.

Parke [12] took a parametric approach to defining the face and representing facial actions. This was an early but impressive approach to computer simulation of the human face, classifying the face by a set of parameters defining size attributes of the facial subsections. Expressions were coded as variations of these parameters, and animations were performed by interpolating along the change in expression. Using this technique, to encode subtle interactions of the face or more complex actions such as bulge creation would require increasingly larger and larger sets of parameters. The system does handle jaw motions remarkably well; this is due to the data being "hard-wired" in -- any effects of an action such as this must be previously known (such as the effect of a widely open jaw on certain cheek and lip actions), so in this respect the system loses some generality. (It should be noted that we do not currently handle jaw actions -- to process these "naturally" is a very difficult problem.)

The last representation considered involves the complete low-level simulation of the face. It is possible to simulate points on the skin, muscles, and bone by a set of interconnected 3-dimensional networks of points, using arcs between selected points to signify relations.

The outermost level or surface is the skin. The skin can be viewed as a continuous 2-dimensional surface, warped and distorted around an ovoid. To simulate the skin, points with 3-dimensional coordinates are selected for the surface. As in Parke's model, points are more concentrated around the most detailed sections -- the eyes and the mouth (most expressions use mainly these portions of the face). Arcs connect a point to all of its "close" neighbors -- that is, any neighbors such that motion of the skin point would affect the position of the neighboring point.

At the innermost level there is the bone structure. This is fairly simple to implement -- it is an inflexible surface at some distance below the skin (ignoring the jaw and its related motions).

Between the bone and skin layers, and spanning the space between them, are the muscles. A muscle is a group of points (muscle fibers) with an arc stretching from each fiber-point to a point on the underlying bone, and another stretching from the fiber-point to one or more skin points.

It is necessary to keep various pieces of relevant information on the arcs. As the different sections vary in their elasticity (resistance to change in position or arclength), information such as length parameters and elasticity information must be stored on the arc.

The basic action performable on the network is the application of a force, or tension, to a select portion of the net. From here, the force propagates outward, affecting more and more distant sections of the model. Since the tensions are integral to the manipulation of the face, these networks are referred to as "Tension Nets".

The Design of the System

The Selected Structures: When we decided to select a pair of representations, we first chose the representation of the face, and then chose the main representation which would work best with the facial representation. Since all of the motion schemes could be used in some fashion for the image analysis (with appropriate extensions where necessary), it seemed to make sense to choose the motion scheme which worked best with the most accurate and usable facial representation.

Of the three examined facial structures, the one which can yield the most usable representation of the face is the method of tension nets. This is based primarily on the fact that it is a "naturally" based system -- the "handles" on the represented facial structures are exactly the same as the motivators of facial actions -- the muscles. Any nuances of the face should naturally fall out as a result of the simulation. This, as we later demonstrated, is indeed the case.

When comparing the various notation schemes to the tension net representation for compatibility, one notation stands out as the clear and obvious choice -- FACS. The output of FACS is a list of currently performed Action Units. Each AU corresponds directly to one or a few muscle contractions or relaxations, precisely the input the tension net scheme requires. Thus, FACS-tension nets offer the following features:

* Any performable action can be simulated. If an action can be notated, the notation yields a simulation. The notation was designed to be able to handle any combination of actions.

* Naturally based system. Analysis and simulation of the face are based on the actual structures of the face. Therefore, any constraints or peculiarities of the real face should appear within the system.

  Close relation between the notations. This insures a simple translation between notation schemes.

* Close relation between the causes of the actions and their simulation. If a muscle causes an action, then in the simulation, the simulated muscles will cause a similar action.

* Independence of FACS from any particular face. The notational scheme was defined to key on changes in features when compared to the subject's base (neutral) face, a well-defined item.

* Uniqueness of decomposition in FACS. Since each facial expression has a unique representation in FACS, this insures reproducability of processing, as well as cross checking of processing with human

notators.

* Efficiency of representation. Two very compact notations signifying the action are available -- the list of AUs being performed, and the status of the muscles of the face.

* Extensibility of the theory. The general theory of simulating an object by the causes of its changes can be extended to cover other nonrigid objects under different circumstances. (See the section Unsolved Problems.)

The System's Subprocesses: The processing performed by each segment can now be further described. The Camera Processor (not yet implemented) would scan the input image for certain facial features. From this list of features, the list of AUs being performed can be determined. This is a nontrivial but well defined process, as AUs are not exclusionary in their interactions. AUs can combine to produce secondary features, they may mask one another, or the performance of one may exclude performance of another. Examples of these processes are (respectively) certain sets of lip actions that modify the cheek shape when performed together (when the cheek would not be changed by the separate actions); eyebrow raising in a fashion that makes detection of certain eyelid actions impossible; and "contradictory" actions such as eyelid droop and eyelid open.

The AU Parser takes the list of AUs and their interactions and looks up the muscle changes for each one. In certain cases, muscle actions have to be combined to give an overall picture of the muscle status of the face.

This list of muscle changes is then passed to the facial Simulator, which performs any necessary tensing or relaxing of muscles. Since the status of the face is saved between camera frames, only actual changes in muscle actions are used as input to this stage. This also insures that parallel actions will be performed in parallel (at the same camera frame), while sequential actions are performed sequentially. This is important, since in many cases the effect of an AU depends on what has been previously performed (the previous structure) and on what is simultaneously being performed (allowing for offsetting actions over sections of the face).

The Data Structures: We constructed the representation of the face in a bottom-up fashion, by starting at the lowest level (a single location on the skin), and building successively higher and higher levels of interaction until sufficiently high level constructs (such as the AU, the basic manipulative unit of FACS) were obtained. These data structures and interactions are described in more detail in Platt[13].

The simplest structure is the point, a 3D location. It represents a single coordinate either on the skin, or as part of the underlying structure (muscle or bone) -- it is a tissue location. Arcs are used to connect points -- an arc exists between two points p1 and p2 iff p1 adj p2. Arcs contain information on the nature of the matter between the two points -- basically, pertaining to the elasticity (spring constant) of the area between p1 and p2. The basic action of force application to a point treats the face as a network of springs, joined at the points. When a force df is applied to a point p, the change in location is computed by

$$d\underline{l} = d\underline{f} \;/\; k$$

where k is the sum of the spring constants at the point.

The simplest logical structure for organizing force applications is the muscle fiber. It consists of a muscle point, to which the force will be applied, a bone point, forming a solid base for the contraction, and one or more skin points, the head(s) of the fiber. Each fiber consists of the muscle point and the magnitude of the force to apply at that point. A simple muscle fiber is shown in figure 2.

From here, fibers are collected into muscles, the basic unit of which contractions are performed. So when a contraction is being specified, each individual fiber has its distinct force applied along it, all performed in parallel. (This is further described in the section Algorithms.)

The highest structure needed in the simulation is that of the AU. Since FACS defines AUs in terms of muscles, each AU consists of of a list of muscles and relative magnitudes. (For example, an AU pulling a muscle at half its maximal force would have a magnitude of 0.5.) It should be noted here that it is at this level of notation that translation from FACS notation to lower-level notation takes place -- a FACS action specification is translated into a list of muscles and forces, and from there to individual point-force combinations.

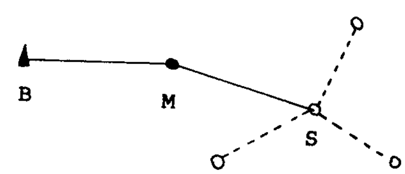

Fig. 2: Muscle Fiber

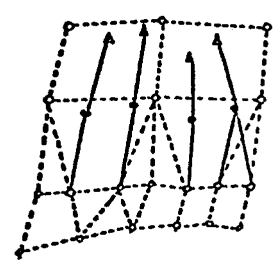

Fig. 3: Muscle

Algorithms: In this section we shall describe the basic techniques used to apply a simple muscle pull to the current facial structure. The algorithm for multiple simultaneous pulls is a simple extension of this, and is also described. The contraction algorithm is an iterative simulation of what is really a continuous event. When sufficient resolution is used for each pull, this proves to be a sufficient simulation.

When a simple fiber contracts, a force is applied to a muscle point in the direction of its tail (bone point). This force causes a displacement of the muscle point; the amount of the displacement varies with the elasticity of the flesh at that point. The force is then reflected along all arcs adjacent to the point; these reflected forces are then applied to their corresponding adjacent points. In this fashion, a force is propagated out from the initiating point across the face.

In actuality, we are dealing with sets of points and forces at all times. To simulate a single muscle or a set of muscles contracting, we only consider the sets of all the fibers (and their applied forces) from ALL the muscles currently contracting.

As previously mentioned, a muscle consists of a set of fibers and their appropriate force-magnitudes. To contract a muscle we apply a force of the prescribed magnitude, along a unit vector stretching from bone to skin of the fiber body, to the muscle (fiber) point itself. In a similar fashion, an AU can be decomposed into a list of muscles, each with an appropriately calculated portion of the magnitude (partial contraction), to be further decomposed from muscle-magnitudes into point-forces.

Animation: Applying a set of actions to a face is not done with one single application of the contraction algorithm. Rather, to help distribute effects of simultaneous pulls over intermediate areas of the face, the contractions are divided into n contractions each of force 1/n. This is analogous to Euler's method for solution of differential equations -- smaller and smaller step sizes yield results more and more interaction, at a computational cost.

Animating a facial action is quite simple once this iteration is considered. The intermediate results are all valid expression complexes, and, as such, the intermediate results are displayed as separate frames. Since partial contractions of muscles are performable, temporally offset animations are easily specifiable by splitting up the AU's into the overlapping and non-overlapping contraction times. This also means that expressions can be performed at different rates even if occurring simultaneously.

Results: A system was implemented using the above structures to perform (animate) facial actions on either a muscle or AU level. The output of this program was fed into a display program, resulting in an image of the arc-lines on the Vector General 3404 (a vector-oriented graphic display device). This image can be rotated about the X-, Y- and Z-axes for further examination. A series of AU contractions were performed to demonstrate the effects of various combinations of actions on the forehead region of the face. The contractions shown here are restricted to those of the upper face. Figures 4a, 4b, and 4e are different views of the neutral face. Muscles are displayed in 4a and 4b; 4e (and most of the rest of the figures) have the muscles omitted for reasons of clarity.

The AUs defined and demonstrated below are:

AU 1: Inner brow raiser. Raises the inner brow by contracting the inner frontalis muscle. The notation "R1" (figures 4c, 4d) is interpreted as the performance of AU 1 only on the right side.

AU 2: Outer brow raiser. Raises the outer brow by contraction of the outer frontalis muscle (figure 4g, 4i). Without training in the performance of AU 2, many people will also contract AU 4, as in figure 4j.

AU 4: Brow Lowerer. Lowers the inner brow by the contraction of the corrigator and pyramidus nasi muscles. It also may create wrinkles at the root of the nose and/or bulges running from mid- to inner- eyebrow. These instances of feature creation are demonstrated in figures 4m (neutral face) and 4n (AU 4 contracted) in the vicinity of the glabella.

AU 6: Cheek Raiser and Lid Compressor. Raises the cheek and narrows the eye opening by contracting the outer portion of the orbicularis ocularis. Figure 4k shows a one-sided contraction (AU R6), performed on the model's right side. Figure 4l demonstrates the combination 4+6, raising the cheeks, lowering the brow, and causing a buckling (creation of a fold) around the inner portion of the upper eye opening.

AU 7: Lid Tightener. Pulls upper and lower eyelids towards the inner corner by contracting the inner orbicularis ocularis. As our model does not currently have eyelids, the effects of this AU are minimal and have not been demonstrated.

Camera Recognition: Currently, research is being initiated to complete this system by constructing an image analysis module. We believe that although difficult, this module is realizable, as there are several constraints upon the (general) image analysis problem which we can apply.

First, the face has a defined structure. Initially, we plan to deal with aligned images (actor performing actions face forward into camera; predefined images, etc.), thus the general features of the face will appear in constrained locations. Interframe alignment and location will be simplified by this constraint. In addition, this further constrains the location of static (eyes, etc.) and dynamic (wrinkles, mass motions) features.

There exists a "base" or neutral image to start with. This is the expressionless actor, and provides a base against which later images can be compared for feature differences.

Finally, the FACS features are predesignated: we know where to look and what to look for. Once features are found, the translation to FACS is a

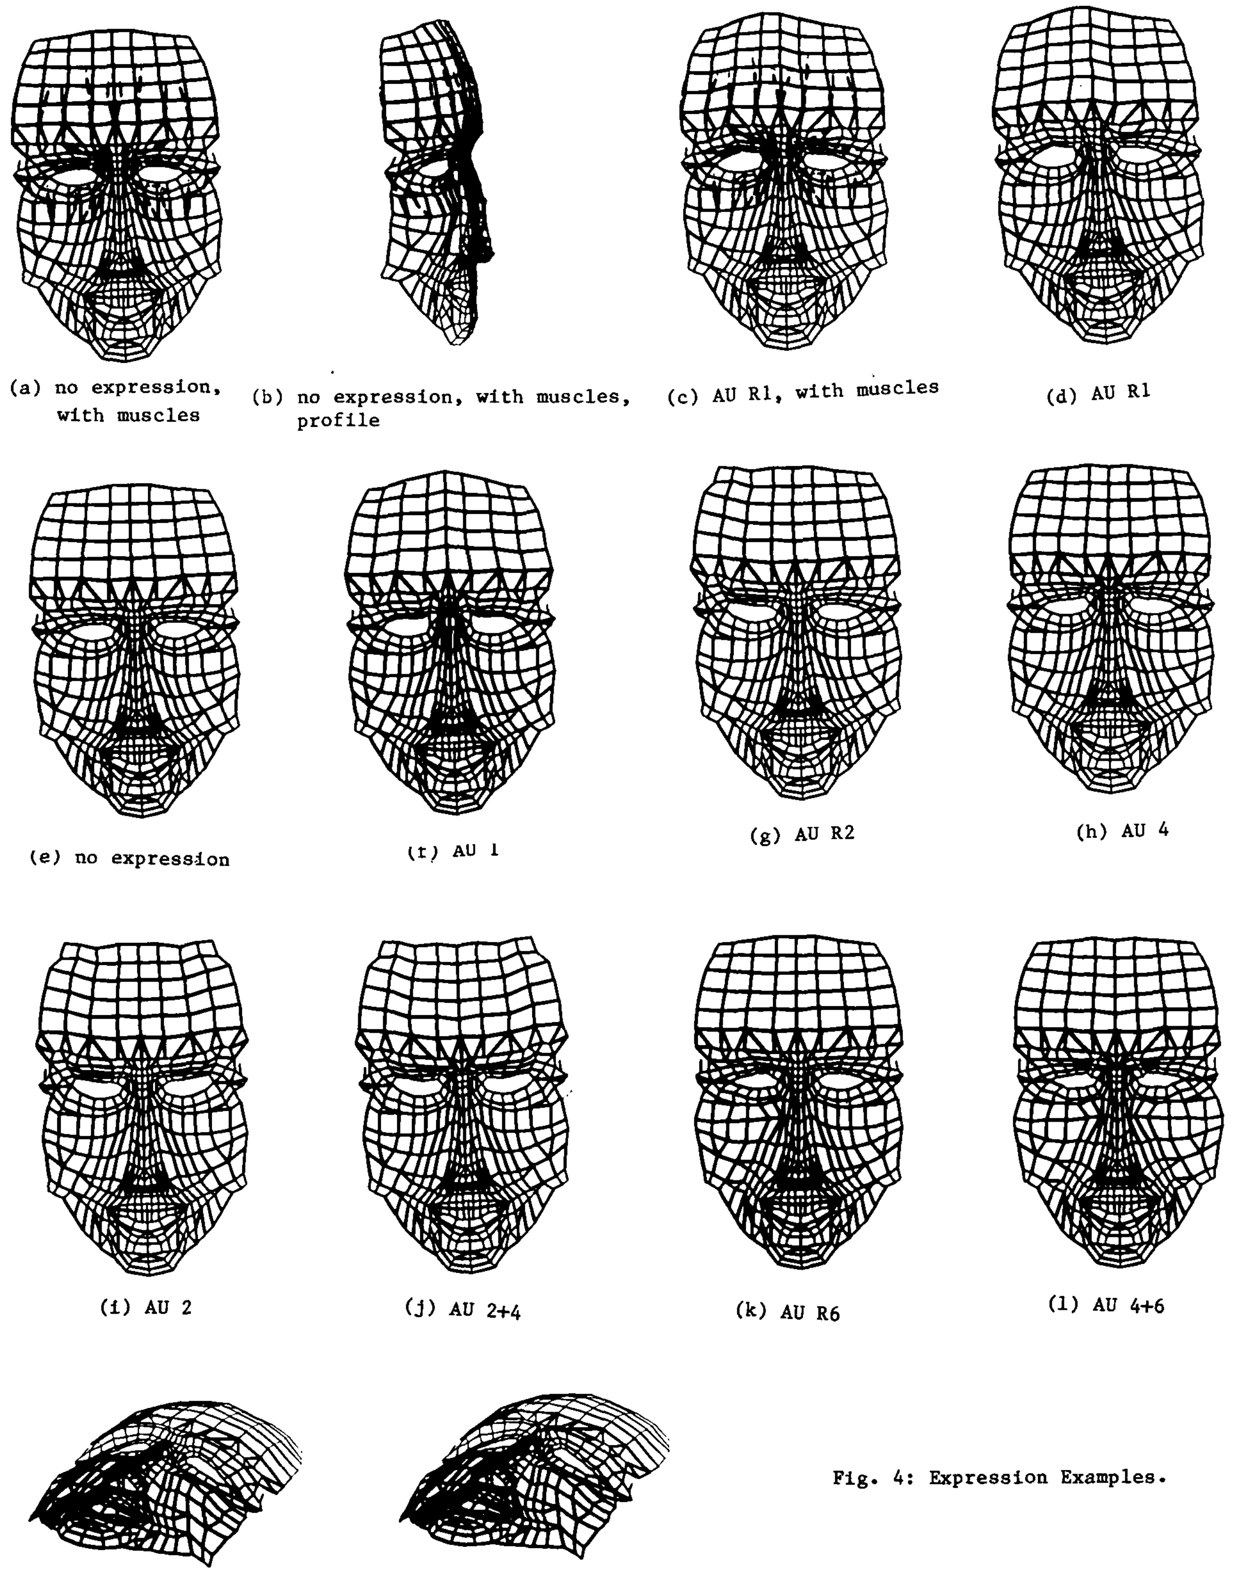

(a) no expression, with muscles

(b) no expression, with muscles, profile

(c) AU R1, with muscles

(d) AU R1

(e) no expression

(f) AU 1

(g) AU R2

(h) AU 4

(i) AU 2

(j) AU 2+4

(k) AU R6

(l) AU 4+6

(m) no expression, chin view

(n) AU 4, chin view

Fig. 4: Expression Examples.

well-defined, although somewhat complex, process.

Some Assorted Problems: Initially, several problems appeared which we thought would have to be handled by special case processing. This was later shown not to be the case, as all desired results fell out naturally as a result of the simulation. This provided further evidence that our system was similar enough to FACS and the (real) face to be confidently used for general animation and simulation.

In the design and learning of FACS, some time is spent on the different forms of AU interactions. The processing of alternating (either of 2 AUs can be alternately scored), parallel (additional side effects due to several closely related AUs), and noninteracting seemed relatively easy, however, the problem of AU masking was still left. An AU will mask another AU when previous performance of the first AU makes the performance of a second AU undetectable. Essentially, the signs characterizing the second AU are either present, or the skin conditions present after the first AU is performed (bulges, wrinkles, skin stretching, etc.) are such that any additions of the second AU are undetectable. An example of this is first smiling, then raising one's cheeks, as opposed to just performing the cheek raiser. The cheek raiser naturally pulls up the corner of the mouth. If a smile has already been performed, the mouth corners are already raised -- the the cheek raiser does not cause an additional raising. This problem does not surface in our model as a result of the way we process skin elasticities. If one AU is sufficiently strong, the skin will be stretched to the point where any additional force will result in no change -- F/k will approach zero as K becomes very large.

A second problem involved the creation of bulges, furrows, and wrinkles in the skin. These three phenomena are basically caused by the same state -- two forces press the points of the skin towards each other, causing a buckling action. Again, our model handles this well -- they will be created if a point is pushed over another point, since the reflected force will approach zero as the angle of reflection (being increased as the point approaches) approaches a 90 degree angle.

Unsolved Problems: At the current time, there are several topics involving the areas this research involves that are suggested by the work, but have not been approached.

First, the system is not ideal with respect to muscles following the flow of a bone sheet, such as the area at the junction of an eye socket and the brow above it. The sudden change of direction will constrain muscle actions in a manner not handled well by the model as currently implemented -- the flesh should flow over the bone, but not through it. At present, the bone is just represented as tie-points for the muscle fibers; there is no concept of a solid mass to be dealt with. The simplest way to code in these constraints would be to increase the complexity of the muscle fiber, encoding additional muscle action constraints, changing the muscle fiber from as represented in figure 3 to that of figure 5.

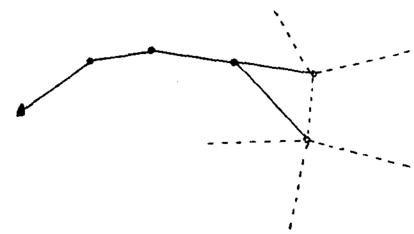

Fig. 5: Improved Muscle

The additional Muscle points are used to define the flow of action of the muscle, while the intermediate Bone points can constrain the vector actions of their related muscle points, thereby forcing the muscle over a bone sheet.

Also, jaw actions are not currently handled by this system. It is believed that they are a clear extension of the structures involved, by treating the bone as a "special case" of action, allowing the muscle actions to pivot the bone around a set of axes. This action would, of course, be propagated through the bone to other connected muscles, causing a stretching action along the skin, and rendering other actions unperformable.

These two problems are combined when dealing with cartilaginous areas, such as the nose. Here, muscles flow over and around the cartilage segments; in addition, they can force the cartilage to move, resulting in a semi-rigid action.

Actions of the cheeks, such as puffing and sucking, were also not handled. To process these actions would require a more complex model of the face and head, involving fluid (air) filled chambers, the actions of the pressures within, and the effects of the bone/stretched skin upon the shapes of the chambers.

Also not investigated were totally nonrigid objects (the facial actions rely on the underlying bone structure for a base for the actions). An example of this within the face is the tongue and its related actions and interactions with the lips and cheeks (the contact problem), as well as pure tongue motions. Another interesting example would be a total heart simulation, enabling the testing of various scenarios of defects/changes of the structure, and their effects on the performance. This would also involve actions of pressure chambers, as the blood filled the various sections of the heart.

Conclusions: We have introduced a model of the human face, utilizing a theory which can be extended to simulate and animate any nonrigid objects whose performable actions are classifiable by primitive interacting units of motion. A simulator for this has been implemented, and overlaid with a notation for actions of the face. Some examples of the output from the system were presented, and possible extensions and other applications were suggested.

References

[1] Badler, Norman I., J. O'Rourke, S. Platt, and M.A. Morris. Human Movement Understanding: A Variety of Perspectives. Proceedings of the First Annual National Conference of Artificial Intelligence, 1980, pp. 53-55.

[2] Badler, Norman I., Joseph O'Rourke, Stephen Smoliar, and Lynne Webber. The Simulation of Human Movement by Computer, Movement Project Report No. 14, Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA. September 1978.

[3] Badler, Norman I., and Stephen Smoliar. Digital Representation of Human Movements. ACM Computing Surveys 11, March 1979, pp. 19-38.

[4] Baker, Charolette. Regulators and Turn-Taking in American Sign Language Discourse, from New Perspectives on American Sign Language, Academic Press, 1977.

[5] Baker, Charolette, and Carol A. Padden. Focusing on the Nonmanual Components of American Sign Language, from Understanding Language through Sign Language Research, Academic Press, 1978.

[6] Birdwhistell, Ray L. Kinesics and Context, University of Pennsylvania Press, Philadelphia, PA. 1970

[7] Calvert, T. W, J. Chapman, and A. Patla. The Integration of Subjective and Objective Data in the Animation of Human Movement. Computer Graphics, vol. 14, no. 3, July 1980.

[8] Ekman, P. and W. Friesen. Facial Action Coding System, Consulting Psychologists Press, Palo Alto, CA. 1978.

[9] Hackathorn, Ronald J. ANIMA III -- A 3-D Color Animation System. Computer Graphics, vol. 11 no. 2, Summer 1977

[10] Hutchinson, Ann. Labanotation. New York: Theatre Arts Books. 1970.

[11] O'Rourke, Joseph P. and N. I. Badler. Model-Based Image Analysis of Human Motion Using Constraint Propagation. IEEE PAMI, 2(6), Nov. 1980, pp. 522-536.

[12] Parke, Frederic I, A Parametric Model for Human Faces. Ph. D. dissertation, University of Utah, Department of Computer Science, 1974.

[13] Platt, Stephen M, A Computer Model for the Simulation of Human Faces. MSE thesis, Department of Computer and Information Science, University of Pennsylvania, 1980

[14] Sutton, Valerie. Sutton Movement Shorthand: Writing Tool for Research, The Movement Shorthand Society Press, Newport Beach, CA. 1978.