The Challenge of Human Figure Animation

Tom Calvert

Centre for Systems Science
Simon Fraser University
Burnaby, B. C. V5A 1S6
(tom@lccr.sfu.cdn)

ABSTRACT

The two major difficulties which must be addressed in producing realistic human figure animation are related to the design of bodies and the specification of movement. Body models can be designed to any desired accuracy for a fixed stance but if the body is to move, a satisfactory way has to be found to automatically adjust the shape of the tissue around the joints. Most animation produced to-day uses kinematic movement specification and requires many hours of work by a skilled animator. Future developments will make use of knowledge of the dynamic properties of bodies and expert systems will automatically define feasible movement patterns. This paper summarizes work going on in this field and describes in more detail several projects at Simon Fraser University which address various aspects of the problem.

RESUME

Les deux principals obstacles à la production d'animation réaliste de figures humaines ont rapport à la création du modèle du corps et à la description du mouvement. Un modèle du corps humain peut être créé à n'importe quel niveau de détail pour une position fixe. Mais si le modèle est animé, il faut trouver un moyen satisfaisant d'ajuster automatiquement la forme des articulations. La plupart des animations aujourd'hui sont produites à partir de descriptions cinématiques du mouvement qui exigent plusieurs heures de travail de la part d'un animateur habile. Les systèmes d'avenir, basés sur la connaissance des propriétés dynamiques du corps humain et sur les systèmes experts, détermineront automatiquement le choix des mouvements possibles. Cet article résume les travaux de recherche sur ce sujet et décrit en plus de détail les projets à l'Université Simon Fraser adressants différents aspects de ce problème.

KEYWORDS: Computer animation, human movement, kinematics, dynamics, expert systems, movement representation, modelling.

THE CHALLENGE

While the animation of a few solid objects is challenging, the animation of human figures is so difficult that it is almost never attempted in production animation. In fact, much of the computer generated animation on television is restricted to "flying logos". The two major difficulties are found in the design of realistic body models and in the specification of realistic movement. While it is quite straightforward to produce reasonable movement on bodies which approximate the human form, we are not very close to producing either bodies or movement which would fool an observer into thinking that they were real. One of the reasons this problem is so difficult is that all humans are sensitized to observe the movement of others very carefully indeed, since body language is an important element of interpersonal communication.

Both in movement specification and body design, the best results are obtained when an animator/designer works for many hours to meticulously define the path or the shape by hand. Thus, it is not surprising that much of the current research concentrates on ways to take advantage of *a priori* knowledge about the problem. One line of research seeks ways to take advantage of the dynamic equations which govern the movement of all bodies and another takes a more general approach to the development of rule based systems to guide the animation. In this paper I will summarize the work going on in this field and describe in more detail several projects at Simon Fraser University which address various aspects of the problem.

THE DESIGN OF BODIES

It is fortunate that the human body is based on an articulated frame, the elements of which are essentially rigid - this is the human skeleton. Thus, most body models are based on a simplified human skeleton, typically with 22-26 segments — an example is shown in Figure 1. This means that the position and configuration of the body in space can be defined by specifying a reference point and the orientation of each of the segments in the articulated structure. In general, the reference point has three dimensions and

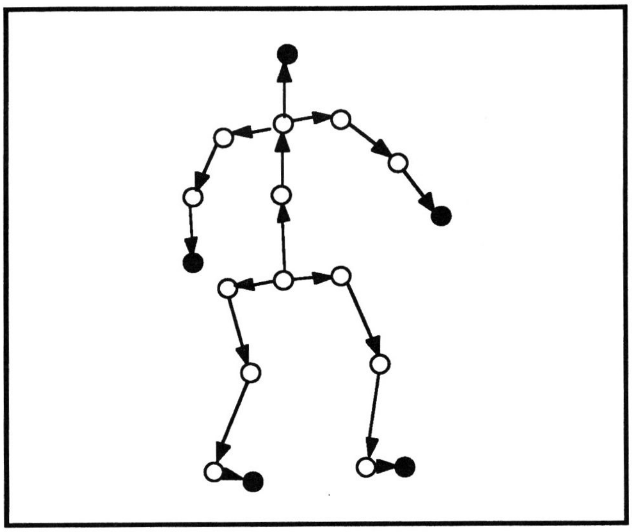

Figure 1. A linked structure representing the parts of the human body

three angles are required to specify the orientation of each segment so we need 69 parameters to fully specify a body with 22 segments. This points out the difficulty of human animation — on the order of 70 parameters need to be specified for each body in each frame, even for a relatively crude figure. If the hands are to have independent fingers or if the face is to be animated, then many more parameters are required.

Because of their simplicity and because of the fact that they do define the body, stick figure displays based on the skeletal model have been widely used. Indeed, David Zeltzer has produced some very nice animation of a full skeleton[1]. Unfortunately it is not completely straightforward to put flesh on the bones. The difficulty is that the tissue surrounding the underlying frame is not rigid and, in particular, it is difficult to find a consistent way to describe the contours of the area around a joint as the body moves through a range of postures. For any fixed posture, standard modelling software can be used to approximate the shape of the body to any desired degree of accuracy. However, for each new posture this will generally need to be readjusted by hand. The difficulty is that the tissue around the joints does not behave in a simple way.

In an attempt to find a simple but robust approach to human body modelling, Korein and Badler [2] developed the "Bubble Man". In this approach the body is built up from the superposition of many hundreds of spheres (Figure 2). This is quite efficient because the sphere is a simple primitive and most importantly, the spheres allow joints to be modelled fairly well for a wide range of postures. However, although the models produced in this way are efficient and robust, there are difficulties in refining them to give a truly realistic look. Another approach based on simple primitives is found in Herbison-Evans' "Sausage Woman"[3]. In this method the body is built up from a small number of ellipsoids (Figure 3). This is also an efficient and robust, but probably not a fruitful route to realism; this method does, however, produce quite good outline figures for a vector display.

As noted above, there are no real technical problems in modelling a body to any desired accuracy; this can be done either with a mesh of polygons or with bicubic patches. The difficulty is that for each new posture the joints must be "fixed up" by hand (e.g. "Tony de Peltrie" [4]). For this reason many examples of human animation avoid the issue in one way or another — often by leaving the joints as gaps (e.g. "*BRILLIANCE*" [5]). Research on solid modelling for CAD/CAM has looked at the related problem of automatically finding blending functions to adjust the shape of a machine part as its shape is changed interactively [6]. This approach is quite promising, but knowledge has to be built in to define the particular way in which the tissue behaves around each joint. Magnenat-Thalmann and Thalmann [7], have described joint dependent local deformation (JLD) operators. These JLD operators control the evolution of surfaces and can be considered as operators on these surfaces. With this approach, a body shape state is a frame-dependent state defined by angles for the frame, the neutral synthetic actor body and the JLD functions for the knee, shoulder, elbow, finger, etc. In a related study, Tony Chung at SFU has investigated approaches to modelling the knee and elbow joints [8]. The approach which he implemented uses stored primitives based on a spline model for the joint shape at a number of angles (say 4). At intermediate angles the shape is obtained by a weighted interpolation.

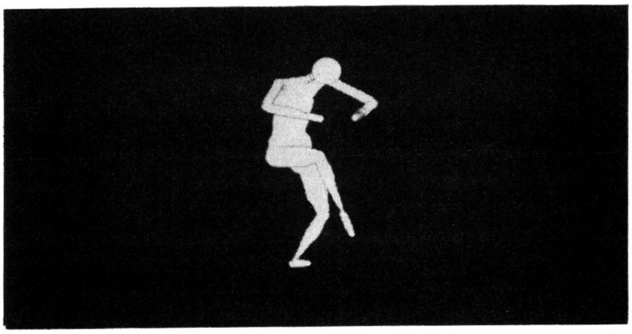

Figure 2. Figure using spheres.

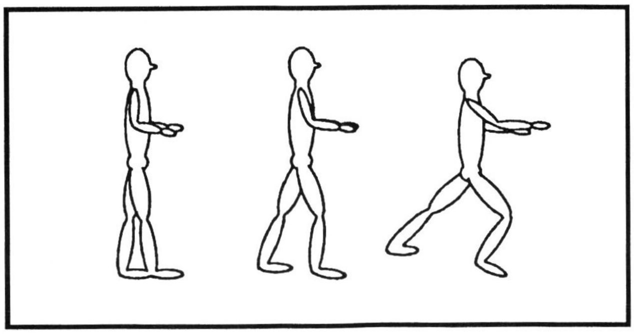

Figure 3. Figure using ellipsoids.

This direction appears promising and we continue to investigate it. However, joints such as the shoulder, which allow movement on all three axes, are particularly challenging.

Modelling the face and capturing accurate facial expressions presents a rather different problem. The movements of the surface are much smaller than around the joints, but they are more subtle and human observers are particularly attuned to the analysis of facial expression. Some animations have used interpolations between expressions digitized from live actors to create a wide range of facial expressions [4]. This can be quite effective, but a more basic approach models changes in the facial surface from muscle movements [9, 10]. In the long run this should give better results.

Relatively little attention has been given to the difficult issue of clothing. While there is current work on the modelling and display of fabric [11], a major difficulty is to find algorithms to describe the way in which the clothing moves relative to the body.

SPECIFICATION OF MOVEMENT

Ideally we would like to develop a flexible and comprehensive computer based language system for human movement. The language system should allow natural graphical and textual inputs at three levels: functional or goal oriented movement, complete co-ordinated movement specification and the lowest level of detailed limb movement specification. The reality of movement specification today is that we have some very well developed kinematic methods and research which points to new approaches which take advantage of dynamic simulation and expert systems.

KINEMATIC METHODS

The animation of an articulated structure requires that a support position be defined in three dimensions and that three angles be specified for each body segment. Using kinematic methods, this can be obtained from live action, from notation or by interactive graphical specification.

Rotoscoping. The most basic approach is to capture the actual movement patterns of a live subject. This can be done by "rotoscoping" which involves digitizing the joint co-ordinates of all body segments from at least two orthogonal views recorded on film or video. This tedious approach is also important in biomechanics research, and there is a continuing interest in automating it, but the pattern recognition problems involved are quite difficult (an excellent example of this approach at its best can be seen in the Robert Able Associates SIGGRAPH Video "The Making of Brilliance" [5]).

Instrumentation. Live movement can also be captured in real-time with special instrumentation. Goniometers provide a relatively inexpensive method and we have experimented with this approach [12]. The major disadvantage is that the measuring equipment inhibits the free movement of the subject. A subject can move more freely with lighter time-multiplexed light emitting diodes attached to the joints; the SELSPOT and WATSMART systems are commercial examples of this rather expensive approach. A promising approach which is currently under development involves a specially instrumented bodysuit which gives the computer real-time analog signals proportional to the angle of each joint. This is an extension of the use of goniometers and the DataGlove [13]. The movement patterns digitized with any of these methods can be normalized for speed and body size and stored in a library of fundamental movement patterns. Much of the commercial animation of human figures has used one of these approaches.

Notation. Movement patterns can also be specified with notation; this approach has been used for recording dance for many decades [14,15,16]. A number of computer based systems have been developed to edit and/or interpret Labanotation [17,18], Benesh notation [19,20] and Eshkol-Wachmann notation scores. Our own experience with Labanotation has shown that it is a viable method to specify animation, but that even with the addition of a macro capability, it is extremely tedious (some might compare it to programming in assembly language) and even dancers do not find it easy to learn. However, this approach has the great advantage that it relies on the animator's conceptualization of the movements required, and it certainly lends itself to the development of complicated scores. A major advantage of this approach is that the resulting score can be edited and changed.

Interactive Specification. A third approach, which has been developed at Simon Fraser [21] and elsewhere [22], involves the interactive specification of body positions in a 3-D graphics environment. In this approach the user is presented with a space-filling vector representation of the human body on the screen of a graphics workstation. The body can be viewed from any angle with a perspective projection under the control of a mouse or equivalent device. The body segments can be picked (again with a mouse or equivalent device), and the orientation of each segment is specified by the user. The end result is directly equivalent to the output from notation but the user has direct visual feedback and finds the adjustments to be reasonably natural and intuitive.

In order to develop a sequence of animation on the computer the animator performs four main steps: the design of a sequence of body positions, the interactive generation of the keyframes, the calculation of intermediate frames by interpolation and a motion test of the results. Each keyframe position is determined interactively on the screen of an IRIS 2400. At the start of the interactive process the user is presented with a body in a standard initial position. The body is represented by a vector display where each body segment is modelled with a four-sided prism. We feel that it is important to use even a crude space filling model to give the user feedback on limb rotation and contact between body segments. Hidden lines are not removed although adjacent body segments are drawn with different colours to aid discrimination. The body parts themselves are positioned individually by a pick and drag procedure.

As the interactive process commences, the user initially selects the foot which is to act as support. Then the body position is built up by orienting the limb segments in turn, moving away from the support point. A body segment is picked using the mouse and its three angles of orientation are adjusted in turn; the user can change the angle of view at will and a digital readout of the angles is given at the bottom of the display. Typically, after the approximate body position has been achieved, the user will iteratively refine the inter-segment angles until the desired result is obtained. Very little typing need be done by the animator while interacting with the system. Instead, software buttons guide the user through the sequence of steps that result in the final animation.

To specify a second frame, the user can either start with a standard position, as with the first, or the first frame can be copied and used as the starting point for the second. Similarly, the third and successive frames are generated. Currently, the user can specify the positions of two figures at a time.

Animation of Multiple Figures. Animating a single figure involves a significant number of parameters for each key-frame. With multiple figures, the problem becomes difficult for the human animator. In an attempt to study how scenes with multiple figures can be composed, we have been working with a dance choreographer, Catherine Lee. We have developed a prototype system for outlining a sequence for multiple figure animation, where the paradigm is parallel to that of the outline processors for text[23].

The animator first sees a screen which is divided into a number of display areas (windows). Down the right side there is a menu of figures in standard stances; up to 14 figures in standing, lunging, sitting, kneeling or lying positions can be displayed at any one time. There are three display areas; the two smaller displays at the bottom show a floor plan and a frontal view of the stage. In the middle of the screen, the third, largest display provides a view of the stage which the animator can continuously adjust by using the mouse.

The animator starts by setting up an initial scene. Figures are chosen from the menu of stances and are placed on the floor plan using the mouse. Their facings are then individually adjusted. Different figures are identified with different colours. When the initial scene is satisfactory it is stored and a second one is created; this is repeated for as many scenes as are needed to define this segment of the animation. These scenes are similar to the series of storyboard sketches used in planning a film, but the interactive 3-D workstation allows the animator to zoom-in or zoom-out from the stage and to view it from any angle. Although the system as it now exists is technically rather straightforward, we believe that it will be a useful tool for further study of the compositional process and that much of what is learned will be transferable to film, video and theatre.

DYNAMIC APPROACHES

An obvious approach to the specification of natural movement is to make use of our knowledge of the dynamic properties of physical objects, i. e. they must obey Newton's Laws of motion. Although the simulation of the movement of simple solids in frictionless space is straightforward, the dynamics of an articulated structure with the complexity of the human body is less simple. Problems arise when estimating the dynamic parameters (moments of inertia, centres of mass, joint friction, muscle/ligament elasticity, etc.), and solving a system of dynamic equations turns out to be computationally rather expensive. The biggest challenge with this approach, however, has been to find intuitive ways to specify and control motion, since it requires knowledge of the joint torques and external reaction forces to solve the dynamic equations and produce animation. Joint torques are produced by the human muscles and are generally under voluntary control. For a moving human subject, these joint torques are in a real sense, unknowable, although methods exist to estimate them. In spite of the fact that it is difficult to estimate parameters and torques, this approach is quite attractive since even if the estimates are inaccurate, the kind of movement which is produced will have some "natural" qualities.

Those investigating this approach include Wilhelms [24], Armstrong [25] and Badler [26]. Human figure models have been produced which, for example, react in a natural way when pulled from a sitting to a standing position. It is not so clear that the complete simulation approach will be useful in generating complex voluntary movement.

A HYBRID APPROACH

Certain types of movement, like walking or running, are conceptually well understood, yet a complete dynamic simulation of such actions is infeasible because generating the proper torques for a locomotion cycle is complicated by problems like balance and co-ordination of the legs. On the other hand, kinematic approaches are inflexible and typically produce a "weightless", unrealistic animation. At SFU, Armin Bruderlin has adopted a hybrid kinematic/dynamic approach. Much like Zeltzer [27], he incorporates knowledge of the locomotion cycle into a hierarchical motion control process. At the top level, the user can specify the parameters of locomotion like step length, step frequency or velocity. The system chooses the right gait and determines the duration and final conditions of the impending step. It then triggers and guides the low level dynamic motor programs for the different phases (stance, swing). Because these motor programs are implemented as dynamic equations, (e.g. compound pendulum for the swing leg), a realistic looking motion is generated. Initial results are encouraging. Our experience suggests that this hybrid of dynamic and goal directed dynamic techniques provides a useful method of animating walking and running under a fairly wide range of conditions.

KNOWLEDGE BASED APPROACHES

The notion of using AI and Expert Systems approaches to capture the knowledge and skills of the animator is obvious, but the implementation is not. At one level, it is fairly straightforward to build up a knowledge base of movement characteristics for the body. These will include basic data on the individual being animated (dimensions, moments of inertia, etc.) and such constraints as the range of movement for specific joints. The knowledge base should also include typical movement patterns for locomotion, simple voluntary movements such as grooming habits, shaking hands and characteristics of the individual which determine how the movements are carried out (brisk, lethargic, happy, sad, etc.).

At the highest level, the system should be able to deduce feasible animation for goal oriented tasks. This would include, for example, feasible movements for the 3 or 4 characters on stage in a television situation-comedy, or a feasible animation for a crowd scene with a dozen people walking up and down a sidewalk (note that people follow quite tight rituals for their body movements when they meet strangers, those known slightly or close friends) [28]. This can also build on work such as that of Reynolds on flocking [29].

There has been some research in this direction [30]. Badler [31] has used a constraint base to constrain his dynamic simulations and Drewery and Tsotsos [32] have studied a frame based approach to goal directed animation. The animation produced by the Human Factory system of Magnenat-Thalmann and Thalmann is most impressive[7]. Perhaps the most extensive work on this problem has been by Gary Ridsdale (formally at SFU, now at the University of Utah). In his PhD thesis [33] he describes a system for human figure animation which finds feasible movement paths which take account of the feelings of the characters being animated as well as other environmental factors.

Our goal is to build on this and to develop expert systems which incorporate the "text-book" rules of direction as well as specific information about individual characters in the action. With this information stored in rule bases, the expert system can provide the animator with an initial framework for the action; the animator would then experiment with a number of his/her own alternatives in arriving at one or more sequences.

Another problem which Ridsdale has studied is the placement of figures in the scene. The rule base includes some of the standard rules of direction for the placement of actors on a stage and for their movement as the action proceeds. It also allows the animator to enter personal information about each character. This might take the form of "John hates Mary", "Mr. Jones is the boss of Mr. Smith", "Bill irritates Simon", etc. Then as action proceeds, mainly through ongoing conversation between the characters, feasible movement patterns are predicted. The system is driven by a script which at this point is only used to indicate which character is speaking. Consideration of the meaning of each utterance is well beyond our current plans.

Not surprisingly, it has been found that the different rules of direction and the rules describing the characters are often in conflict. Thus, mechanisms

To illustrate the use of the planning routines in the General Knowledge Library, suppose the following script were to be animated:

Script Fact Collection (includes)

Characters:
(1) Tex (T).
(2) Black Bart (B).
(3) Miss Kitty (K).

Setting:
as in Figure 5.

Relationships:
(1) Tex likes Miss Kitty.
(2) Black Bart desires Miss Kitty.
(3) Miss Kitty desires Tex.
(4) Miss Kitty is afraid of Black Bart.
(5) Black Bart hates Tex.
(6) Tex hates Black Bart.

General Knowledge Library (includes)

(1) Hot stoves are obstacles.
(2) Holes in the floor are obstacles.
(3) Characters tend to avoid hot stoves.
(4) Characters tend to avoid holes in the floor.
(5) People who are sad tend to use gait "shuffle".
(6) People who are angry tend to use gait "stamp".
(7) People who are shy tend to use gait "walk-quickly".

Action (director's annotations in italics)

- Tex says, "Bart, you dirty pole-cat. You'd better be goin'."
- Black Bart grunts and exits outside.
  Black Bart is bowlegged. Black Bart is angry.
- Miss Kitty says, "Going to stay awhile, Tex?"
- Tex says, "No, ma'am. Got some dogies to corral. See ya'll."
- Tex exits outside.
  Tex is shy.
- Miss Kitty sighs and exits to the kitchen.
  Miss Kitty is sad.

Figure 4. Planning Example

must be developed to provide feasible solutions which in some sense minimize these conflicts; obviously there will be no unique results and the animator may wish to apply his judgement to the weighting of different rules. This area of expert systems theory is poorly developed and Ridsdale has had to develop some new approaches. Some of the flavour of this work can be obtained from the Planning Example set out in figures 4, 5 and 6.

Each line of the script represents a snapshot of the action in the scene, and translates directly into a set of key positions for the characters, together with suggestions for the constraints on their actions to the next position. Figure 6 depicts the plan of Black Bart exiting.

General Knowledge library statements are used to activate operator *insert-keypoint* to go around the obstacles. Each of the new candidate keypoints has a tendency pointing away from the stove and the hole, since they are obstacles that should be given a wide berth. The fact that Bart desires Miss Kitty and hates Tex also affects the tendencies from each position. Since the sum-vector for the four key positions ("n") points upstage-left, the upstage path (path 2) around the tendency displacements away from Tex and towards Miss Kitty is chosen. The gait pattern for Black Bart's walk is chosen as a compromise between "bowlegged-gait" and "stamping-gait".

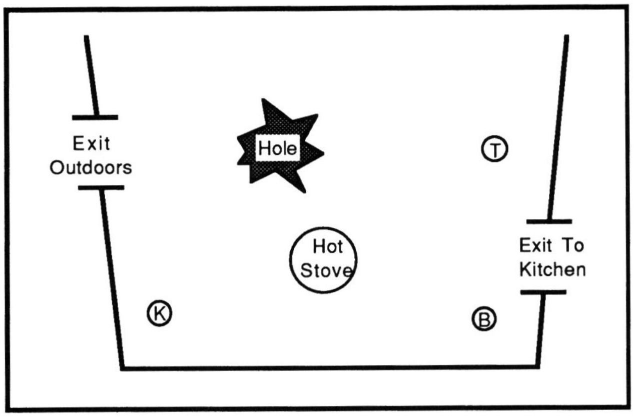

Figure 5. Initial setting

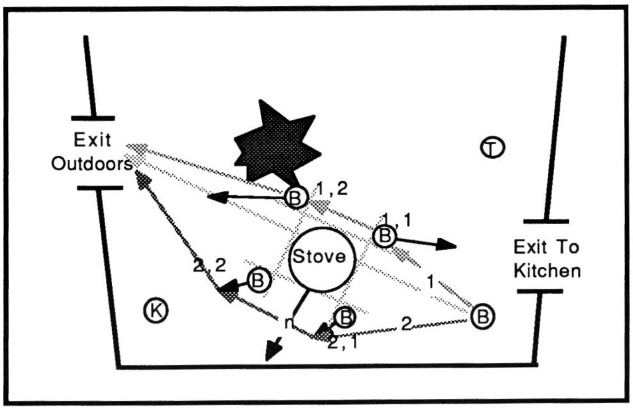

Figure 6. Black Bart exits

THE IDEAL SYSTEM

We believe that an ideal animation system for human figure animation will have the following general characteristics:

1. At the highest level, the system will accept natural language input with about the same information content and complexity found in the script for a film or a play. Presumably the characters to be animated will be identified and their special characteristics will be described where these are considered to be important. There will be standard stage directions (Bill enters stage left, etc.). It is expected that from this input the system will generate a feasible animation. The general nature of such a high level system was discussed in a recent paper [34].

2. At a second, intermediate level, the high level input will generate a detailed script. This script will provide the details of the movement for the figures in terms of gait, style of gesture, etc., as well as the start times and durations for all movements. The important feature of this score is that it is readable to the animator and can be added to or edited. It is in this way that the animator can adjust the feasible animation generated by the system to what is really needed.

3. The third and lowest level comprises detailed movement instructions for each limb segment as a function of time, possibly implemented using hierarchical channel tables. This is also accessible to the animator, but it is assumed that it is seldom changed, except for "fine tuning".

A SYSTEM FOR THE 1990'S

The ideal characteristics outlined above are currently unachievable, but many of the components exist. Ridsdale [33] has shown that the personal characteristics of the figures to be animated can be taken into account. We know that the use of dynamics promises to provide a method of generating locomotion over unpredictable terrain. However, some difficult problems remain. Probably the most difficult will be the animation of the human face. It has been pointed out above that humans critically observe each others movements as a means of communication - the most critical observation is reserved for the face. Another difficulty is clothing. We still do not know how to animate clothing which moves naturally with the body. At least for the meantime, however, we can look forward to very successful animation of crowd scenes where the faces are indistinct and the figures are either nude or clothed in leotards!

ACKNOWLEDGEMENTS

The work described here has been supported in part by grants from the Social Sciences and Humanities Research Council, the Natural Sciences and Engineering Research Council and the Centre for Systems Science at Simon Fraser University. The author gratefully acknowledges the contributions of Severin Gaudet, Gary Ridsdale, Catherine Lee, Tony Chung and Armin Bruderlin.

References

1. D. Zeltzer, "Motor control techniques for figure animation", **IEEE Computer Graphics and Applications**, vol. 2, no. 9, pp. 53—59, November 1982.

2. N. I. Badler, J. O'Rourke and H. Toltzis, "A spherical representation of a human body for visualizing movement", **Proc. IEEE**, vol. 67, pp. 1397—1402, 1979.

3. D. Herbison-Evans, "NUDES-2: A numeric utility displaying ellipsoid solids", **Computer Graphics**, vol. 12, no. 3, pp. 354—356, August 1978.

4. P. Bergeron, "Techniques for animating characters", SIGGRAPH '85 Tutorial Notes, ACM, New York, 1985.

5. Robert Able and Associates, "The making of *BRILLIANCE* ", SIGGRAPH Video Review, issue 20, ACM, New York, 1985.

6. C. Hoffman and J. Hopcroft, "Automatic surface generation in computer-aided design", **The Visual Computer**, vol. 1, pp 92—100, 1985

7. N. Magnenat-Thalmann and D. Thalmann, "The direction of synthetic actors in the film *Rendez-vous à Montréal* ", **IEEE Computer Graphics and Applications**, vol. 7, no. 12, pp. 9—19, December 1987.

8. T. Chung, "An Approach to Human Surface Modelling Using Cardinal Splines", MSc thesis, School of Computing Science, Simon Fraser University, 1987.

9. F. I. Parke, "Parameterized models for facial animation", **IEEE Computer Graphics and Applications**, vol. 2, no. 9, pp. 61—68, November 1982.

10. S. Platt and N. Badler, "Animating facial expressions", **Computer Graphics**, vol. 15, no. 3, pp. 245—252, August 1981.

11. J. Weil, "The Synthesis of cloth objects", **Computer Graphics**, vol. 20, no. 4, pp 49—54, August 1986

12. T. W. Calvert, J. Chapman and A. Patla, "The integration of subjective and objective data in animation of human movement", **Computer Graphics**, vol. 14, no. 3, pp. 198—203, August 1980.

13. DataGlove by VPL Inc., Redwood City, CA.

14. R. Benesh and J. Benesh, **An Introduction to Benesh Dance Notation**, A. C. Black, London 1956.

15. A. Hutchinson, **Labanotation**, Theatre Arts Books, 2nd edition, New York 1960.

16. N. Eshkol and A. Wachmann, **Movement Notation**, Weidenfeld and Nicholson, London 1958.

17. T. W. Calvert, J. Chapman and A. Patla, "Aspects of the kinematic simulation of human movement", **IEEE Computer Graphics and Applications**, vol. 2, no. 9, pp. 41—50, November 1982.

18. N. I. Badler and S. W. Smoliar, "Digital representation of human movement", **Computing Surveys**, vol. 11, pp. 19—38, March 1979.

19. R. Ryman, A. Singh, J. C. Beatty, and K. S. Booth, "A computerized editor for Benesh movement notation", **CORD Dance Research Journal**, vol. 16, pp. 27—34, 1984.

20. G. Politis, "A Computer Graphics Interpreter for Benesh Movement Notation", PhD thesis, Dept. of Computer Science, University of Sydney, 1987.

21. G. Ridsdale, S. Hewitt and T. W. Calvert, "The interactive specification of human animation", **Proc. Graphics Interface '86**, pp. 121—130, 1986.

22. S. Van Baerle, "A case study of flexible figure animation", SIGGRAPH '87 Tutorial Notes, ACM, New York, 1987

23. T. W. Calvert, C. Lee, G. Ridsdale, S. Hewitt and V. Tso, "The interactive composition of scores for dance", **Dance Notation Journal**, 1986.

24. J. Wilhelms and B. A. Barsky, "Using dynamic analysis to animate articulated bodies such as humans and robots", **Proc. Graphics Interface '85**, pp. 97—104, 1985.

25. W. W. Armstrong and M. Green, "The dynamics of articulated rigid bodies for purposes of animation", **Proc. Graphics Interface '85**, pp. 407—416, 1985.

26. N. I. Badler, "Animating human figures: Perspectives and directions", **Proc. Graphics Interface '86**, pp. 115—120, 1986.

27. D. Zeltzer, "Motor control techniques for figure animation", **IEEE Computer Graphics and Applications**, vol. 2, no. 9, pp. 53—59, November, 1982.

28. R. L. Birdwhistell, **Kinesics and Context: Essays on Body Movement Communication,** Univ. of Pennsylvania Press, Philadelphia, 1970.

29. C. Reynolds, "Flocks, herds and schools: a distributed behavioral model", **Computer Graphics**, vol. 21, no. 4, pp. 25—34, July 1987.

30. D. Zeltzer, "Knowledge based animation", **Proc. ACM SIGGRAPH/SIGART workshop on motion**, pp. 187—192, 1983.

31. N. I. Badler et al., "Positioning and animating figures in a task oriented environment", **The Visual Computer**, pp. 212—220, December 1985.

32. K. Drewery and J. Tsotsos, "Goal directed animation using English motion commands", **Proc. Graphics Interface '86**, pp. 131—135, 1986.

33. G. Ridsdale, "The Director's Apprentice: Animating Figures in a Constrained Environment", PhD thesis, School of Computing Science, Simon Fraser University, 1987.

34. T. W. Calvert, "Towards a language for human movement", **Computers and the Humanities**, vol. 20, no. 2, pp. 35—43, 1986.