

Simulating Humans

Computer Graphics Animation and Control

NORMAN I. BADLER, CARY B. PHILLIPS,
BONNIE LYNN WEBBER

*Department of Computer and Information Science
University of Pennsylvania*



New York Oxford
OXFORD UNIVERSITY PRESS
1993

Oxford University Press

Oxford New York Toronto

Delhi Bombay Calcutta Madras Karachi

Kuala Lumpur Singapore Hong Kong Tokyo

Nairobi Dar es Salaam Cape Town

Melbourne Auckland Madrid

and associated companies in

Berlin Ibadan

Copyright © 1993 by Oxford University Press, Inc. and
Norman I. Badler, Cary B. Phillips, Bonnie L. Webber

Published by Oxford University Press, Inc.,
200 Madison Avenue, New York, New York 10016

Oxford is a registered trademark of Oxford University Press

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
electronic, mechanical, photocopying, recording, or otherwise,
without the prior permission of Oxford University Press.

Library of Congress Cataloging-in-Publication Data
Badler, Norman I.

Simulating humans : computer graphics animation and control /

Norman I. Badler, Cary B. Phillips, Bonnie L. Webber.

p. cm. Includes bibliographical references and index.

ISBN 0-19-507359-2

1. Human engineering. 2. Virtual reality.

3. Body, Human—Computer simulation.

4. Human mechanics—Computer simulation.

I. Phillips, Cary B. II. Webber, Bonnie Lynn. III. Title.

TA166.B32 1993

620.8'2—dc20 93-12061

9 8 7 6 5 4 3 2 1

Printed in the United States of America
on acid-free paper

Simulating Humans

To Ginny, Denise, and Mark

Preface

The decade of the 80's saw the dramatic expansion of high performance computer graphics into domains previously able only to flirt with the technology. Among the most dramatic has been the incorporation of real-time interactive manipulation and display for human figures. Though actively pursued by several research groups, the problem of providing a virtual or synthetic human for an engineer or designer already accustomed to Computer-Aided Design techniques was most comprehensively attacked by the Computer Graphics Research Laboratory at the University of Pennsylvania. The breadth of that effort as well as the details of its methodology and software environment are presented in this volume.

This book is intended for human factors engineers requiring current knowledge of how a computer graphics surrogate human can augment their analyses of designed environments. It will also help inform design engineers of the state-of-the-art in human figure modeling, and hence of the human-centered design central to the emergent notion of Concurrent Engineering. Finally, it documents for the computer graphics community a major research effort in the interactive control and motion specification of articulated human figures.

Many people have contributed to the work described in this book, but the textual material derives more or less directly from the efforts of our current and former students and staff: Tarek Alameldin, Francisco Azuola, Breck Baldwin, Welton Becket, Wallace Ching, Paul Diefenbach, Barbara Di Eungenio, Jeffrey Esakov, Christopher Geib, John Granieri, Marc Grosso, Pei-Hwa Ho, Mike Hollick, Moon Jung, Jugal Kalita, Hyeongseok Ko, Eunyoung Koh, Jason Koppel, Michael Kwon, Philip Lee, Libby Levison, Gary Monheit, Michael Moore, Ernest Otani, Susanna Wei, Graham Walters, Michael White, Jianmin Zhao, and Xinmin Zhao. Additional animation help has come from Leanne Hwang, David Haynes, and Brian Stokes. John Granieri and Mike Hollick helped considerably with the photographs and figures.

This work would not have been possible without the generous and often long term support of many organizations and individuals. In particular we would like to acknowledge our many colleagues and friends: Barbara Woolford, Geri Brown, Jim Maida, Abhilash Pandya and the late Linda Orr in the Crew Station Design Section and Mike Greenisen at NASA Johnson Space Center; Ben Cummings, Brenda Thein, Bernie Corona, and Rick Kozycki of the U.S. Army Human Engineering Laboratory at Aberdeen Proving Grounds; James Hartzell, James Larimer, Barry Smith, Mike Prevost, and Chris Neukom of the *A³I* Project in the Aeroflight Dynamics Directorate of NASA Ames Research Center; Steve Paquette of the U. S. Army Natick Laboratory; Jagdish Chandra and David Hislop of the U. S. Army Research Office; the Army Artificial Intelligence Center of Excellence at the University of Pennsylvania and its Director, Aravind Joshi; Art Iverson and Jack Jones of the U.S. Army TACOM; Jill Easterly, Ed Boyle, John Ianni, and Wendy Campbell of the U. S. Air Force Human Resources Directorate at Wright-Patterson Air Force Base; Medhat Korna and Ron Dierker of Systems Exploration, Inc.; Pete Glor

and Joseph Spann of Hughes Missile Systems (formerly General Dynamics, Convair Division); Ruth Maulucci of MOCO Inc.; John McConville, Bruce Bradtmiller, and Bob Beecher of Anthropology Research Project, Inc.; Edmund Khouri of Lockheed Engineering and Management Services; Barb Fecht of Battelle Pacific Northwest Laboratories; Jerry Duncan of Deere and Company; Ed Bellandi of FMC Corp.; Steve Gulasy of Martin-Marietta Denver Aerospace; Joachim Grollman of Siemens Research; Kathleen Robinette of the Armstrong Medical Research Lab at Wright-Patterson Air Force Base; Harry Frisch of NASA Goddard Space Flight Center; Jerry Allen and the folks at Silicon Graphics, Inc.; Jack Scully of Ascension Technology Corp.; the National Science Foundation CISE Grant CDA88-22719 and ILI Grant USE-9152503; and the State of Pennsylvania Benjamin Franklin Partnership. Martin Zaidel contributed valuable \LaTeX help. Finally, the encouragement and patience of Don Jackson at Oxford University Press has been most appreciated.

Norman I. Badler
University of Pennsylvania

Cary B. Phillips
PDI, Sunnyvale

Bonnie L. Webber
University of Pennsylvania

Contents

1	Introduction and Historical Background	1
1.1	Why Make Human Figure Models?	4
1.2	Historical Roots	7
1.3	What is Currently Possible?	11
1.3.1	A Human Model must be Structured Like the Human Skeletal System	12
1.3.2	A Human Model should Move or Respond Like a Human	12
1.3.3	A Human Model should be Sized According to Permissible Human Dimensions	14
1.3.4	A Human Model should have a Human-Like Appearance	15
1.3.5	A Human Model must Exist, Work, Act and React Within a 3D Virtual Environment	15
1.3.6	Use the Computer to Analyze Synthetic Behaviors . . .	16
1.3.7	An Interactive Software Tool must be Designed for Usability	18
1.4	Manipulation, Animation, and Simulation	19
1.5	What Did We Leave Out?	20
2	Body Modeling	23
2.1	Geometric Body Modeling	23
2.1.1	Surface and Boundary Models	23
2.1.2	Volume and CSG Models	25
2.1.3	The Principal Body Models Used	27
2.2	Representing Articulated Figures	28
2.2.1	Background	29
2.2.2	The Terminology of Peabody	30
2.2.3	The Peabody Hierarchy	31
2.2.4	Computing Global Coordinate Transforms	33
2.2.5	Dependent Joints	33
2.3	A Flexible Torso Model	34
2.3.1	Motion of the Spine	36
2.3.2	Input Parameters	37
2.3.3	Spine Target Position	38
2.3.4	Spine Database	38

2.4	Shoulder Complex	39
2.4.1	Primitive Arm Motions	40
2.4.2	Allocation of Elevation and Abduction	41
2.4.3	Implementation of Shoulder Complex	41
2.5	Clothing Models	45
2.5.1	Geometric Modeling of Clothes	46
2.5.2	Draping Model	48
2.6	The Anthropometry Database	49
2.6.1	Anthropometry Issues	49
2.6.2	Implementation of Anthropometric Scaling	50
2.6.3	Joints and Joint Limits	51
2.6.4	Mass	53
2.6.5	Moment of Inertia	53
2.6.6	Strength	54
2.7	The Anthropometry Spreadsheet	54
2.7.1	Interactive Access Anthropometric Database	56
2.7.2	SASS and the Body Hierarchy	57
2.7.3	The Rule System for Segment Scaling	57
2.7.4	Figure Creation	59
2.7.5	Figure Scaling	59
2.8	Strength and Torque Display	60
2.8.1	Goals of Strength Data Display	61
2.8.2	Design of Strength Data Displays	61
3	Spatial Interaction	67
3.1	Direct Manipulation	67
3.1.1	Translation	68
3.1.2	Rotation	68
3.1.3	Integrated Systems	69
3.1.4	The Jack Direct Manipulation Operator	70
3.2	Manipulation with Constraints	75
3.2.1	Postural Control using Constraints	75
3.2.2	Constraints for Inverse Kinematics	77
3.2.3	Features of Constraints	78
3.2.4	Inverse Kinematics and the Center of Mass	78
3.2.5	Interactive Methodology	80
3.3	Inverse Kinematic Positioning	83
3.3.1	Constraints as a Nonlinear Programming Problem	86
3.3.2	Solving the Nonlinear Programming Problem	87
3.3.3	Assembling Multiple Constraints	91
3.3.4	Stiffness of Individual Degrees of Freedom	93
3.3.5	An Example	93
3.4	Reachable Spaces	94
3.4.1	Workspace Point Computation Module	96
3.4.2	Workspace Visualization	97
3.4.3	Criteria Selection	98

4 Behavioral Control	101
4.1 An Interactive System for Postural Control	102
4.1.1 Behavioral Parameters	103
4.1.2 Passive Behaviors	109
4.1.3 Active Behaviors	114
4.2 Interactive Manipulation With Behaviors	116
4.2.1 The Feet	117
4.2.2 The Center of Mass and Balance	117
4.2.3 The Torso	120
4.2.4 The Pelvis	123
4.2.5 The Head and Eyes	123
4.2.6 The Arms	123
4.2.7 The Hands and Grasping	126
4.3 The Animation Interface	126
4.4 Human Figure Motions	128
4.4.1 Controlling Behaviors Over Time	129
4.4.2 The Center of Mass	129
4.4.3 The Pelvis	130
4.4.4 The Torso	130
4.4.5 The Feet	130
4.4.6 Moving the Heels	131
4.4.7 The Arms	132
4.4.8 The Hands	132
4.5 Virtual Human Control	132
5 Simulation with Societies of Behaviors	137
5.1 Forward Simulation with Behaviors	139
5.1.1 The Simulation Model	141
5.1.2 The Physical Execution Environment	142
5.1.3 Networks of Behaviors and Events	144
5.1.4 Interaction with Other Models	145
5.1.5 The Simulator	147
5.1.6 Implemented Behaviors	149
5.1.7 Simple human motion control	150
5.2 Locomotion	150
5.2.1 Kinematic Control	151
5.2.2 Dynamic Control	152
5.2.3 Curved Path Walking	154
5.2.4 Examples	159
5.3 Strength Guided Motion	161
5.3.1 Motion from Dynamics Simulation	161
5.3.2 Incorporating Strength and Comfort into Motion	163
5.3.3 Motion Control	164
5.3.4 Motion Strategies	167
5.3.5 Selecting the Active Constraints	169
5.3.6 Strength Guided Motion Examples	170

5.3.7	Evaluation of this Approach	173
5.3.8	Performance Graphs	173
5.3.9	Coordinated Motion	174
5.4	Collision-Free Path and Motion Planning	180
5.4.1	Robotics Background	180
5.4.2	Using Cspace Groups	181
5.4.3	The Basic Algorithm	182
5.4.4	The Sequential Algorithm	183
5.4.5	The Control Algorithm	185
5.4.6	The Planar Algorithm	186
5.4.7	Resolving Conflicts between Different Branches	186
5.4.8	Playing Back the Free Path	187
5.4.9	Incorporating Strength Factors into the Planned Motion	189
5.4.10	Examples	190
5.4.11	Completeness and Complexity	191
5.5	Posture Planning	192
5.5.1	Functionally Relevant High-level Control Parameters	196
5.5.2	Motions and Primitive Motions	197
5.5.3	Motion Dependencies	197
5.5.4	The Control Structure of Posture Planning	199
5.5.5	An Example of Posture Planning	200
6	Task-Level Specifications	207
6.1	Performing Simple Commands	208
6.1.1	Task Environment	208
6.1.2	Linking Language and Motion Generation	209
6.1.3	Specifying Goals	209
6.1.4	The Knowledge Base	210
6.1.5	The Geometric Database	211
6.1.6	Creating an Animation	211
6.1.7	Default Timing Constructs	212
6.2	Language Terms for Motion and Space	214
6.2.1	Simple Commands	214
6.2.2	Representational Formalism	215
6.2.3	Sample Verb and Preposition Specifications	217
6.2.4	Processing a sentence	219
6.2.5	Summary	221
6.3	Task-Level Simulation	222
6.3.1	Programming Environment	223
6.3.2	Task-actions	224
6.3.3	Motivating Some Task-Actions	225
6.3.4	Domain-specific task-actions	226
6.3.5	Issues	228
6.3.6	Summary	231
6.4	A Model for Instruction Understanding	231

7	Epilogue	243
7.1	A Roadmap Toward the Future	244
7.1.1	Interactive Human Models	245
7.1.2	Reasonable Biomechanical Properties	245
7.1.3	Human-like Behaviors	245
7.1.4	Simulated Humans as Virtual Agents	246
7.1.5	Task Guidance through Instructions	246
7.1.6	Natural Manual Interfaces and Virtual Reality	246
7.1.7	Generating Text, Voice-over, and Spoken Explication for Animation	247
7.1.8	Coordinating Multiple Agents	247
7.2	Conclusion	248
	Bibliography	249
	Index	267

Chapter 1

Introduction and Historical Background

People are all around us. They inhabit our home, workplace, entertainment, and environment. Their presence and actions are noted or ignored, enjoyed or disdained, analyzed or prescribed. The very ubiquitousness of other people in our lives poses a tantalizing challenge to the computational modeler: people are at once the most common object of interest and yet the most structurally complex. Their everyday movements are amazingly fluid yet demanding to reproduce, with actions driven not just mechanically by muscles and bones but also cognitively by beliefs and intentions. Our motor systems manage to learn how to make us move without leaving us the burden or pleasure of knowing how we did it. Likewise we learn how to describe the actions and behaviors of others without consciously struggling with the processes of perception, recognition, and language.

A famous Computer Scientist, Alan Turing, once proposed a test to determine if a computational agent is intelligent [Tur63]. In the Turing Test, a subject communicates with two agents, one human and one computer, through a keyboard which effectively restricts interaction to language. The subject attempts to determine which agent is which by posing questions to both of them and guessing their identities based on the “intelligence” of their answers. No physical manifestation or image of either agent is allowed as the process seeks to establish abstract “intellectual behavior,” thinking, and reasoning. Although the Turing Test has stood as the basis for computational intelligence since 1963, it clearly omits any potential to evaluate physical actions, behavior, or appearance.

Later, Edward Feigenbaum proposed a generalized definition that included action: “Intelligent action is an act or decision that is goal-oriented, arrived at by an understandable chain of symbolic analysis and reasoning steps, and is one in which knowledge of the world informs and guides the reasoning.” [Bod77]. We can imagine an analogous “Turing Test” that would have the

subject watching the behaviors of two agents, one human and one synthetic, while trying to determine at a better than chance level which is which. Human movement enjoys a universality and complexity that would definitely challenge an animated figure in this test: if a computer-synthesized figure looks, moves, and acts like a real person, are we going to believe that it is real? On the surface the question almost seems silly, since we would rather not allow ourselves to be fooled. In fact, however, the question is moot though the premises are slightly different: cartoon characters are hardly “real,” yet we watch them and properly interpret their actions and motions in the evolving context of a story. Moreover, they are not “realistic” in the physical sense – no one expects to see a manifest Mickey Mouse walking down the street. Nor do cartoons even move like people – they squash and stretch and perform all sorts of actions that we would never want to do. But somehow our perceptions often make these characters *believable*: they appear to act in a goal-directed way because their human animators have imbued them with physical “intelligence” and behaviors that apparently cause them to chase enemies, bounce off walls, and talk to one another. Of course, these ends are achieved by the skillful weaving of a story into the crafted images of a character. Perhaps surprisingly, the mechanisms by which motion, behavior, and emotion are encoded into cartoons is *not* by building synthetic models of little creatures with muscles and nerves. The requisite animator skills do not come easily; even in the cartoon world refinements to the art and technique took much work, time, and study [TJ81]. Creating such movements automatically in response to real-time interactive queries posed by the subject in our hypothetical experiment does not make the problem any easier. Even Turing, however, admitted that the intelligence sought in his original test did not require the computational *process* of thinking to be identical to that of the human: the external manifestation in a plausible and reasonable answer was all that mattered.

So why are we willing to assimilate the truly artificial reality of cartoons – characters created and moved entirely unlike “real” people – yet be skeptical of more human-like forms? This question holds the key to our physical Turing Test: as the appearance of a character becomes more human, our perceptual apparatus demands motion qualities and behaviors which sympathize with our expectations. As a cartoon character takes on a human form, the only currently viable method for accurate motion is the recording of a real actor and the tracing or transfer (“rotoscoping”) of that motion into the animation. Needless to say, this is not particularly satisfying to the modeler: the motion and actor must exist prior to the synthesized result. Even if we recorded thousands of individual motions and retrieved them through some kind of indexed video, we would still lack the freshness, variability, and adaptability of humans to live, work, and play in an infinite variety of settings.

If synthetic human motion is to be produced without the benefit of prior “real” execution and still have a shot at passing the physical Turing Test, then models must carefully balance structure, shape, and motion in a compatible package. If the models are highly simplified or stylized, cartoons or caricatures will be the dominant perception; if they look like humans, then they will be

expected to behave like them. How to accomplish this without a real actor showing the way is the challenge addressed here.

Present technology can approach human appearance and motion through computer graphics modeling and three-dimensional animation, but there is considerable distance to go before purely synthesized figures trick our senses. A number of promising research routes can be explored and many are taking us a considerable way toward that ultimate goal. By properly delimiting the scope and application of human models, we can move forward, not to replace humans, but to substitute adequate computational surrogates in various situations otherwise unsafe, impossible, or too expensive for the real thing.

The goals we set in this study are realistic but no less ambitious than the physical Turing Test: we seek to build computational models of human-like figures which, though they may not trick our senses into believing they are alive, nonetheless manifest animacy and convincing behavior. Towards this end, we

- Create an interactive computer graphics human model.
- Endow it with reasonable biomechanical properties.
- Provide it with “human-like” behaviors.
- Use this simulated figure as an agent to effect changes in its world.
- Describe and guide its tasks through natural language instructions.

There are presently no perfect solutions to any of these problems, but significant advances have enabled the consideration of the suite of goals under uniform and consistent assumptions. Ultimately, we should be able to give our surrogate human directions that, in conjunction with suitable symbolic reasoning processes, make it appear to behave in a natural, appropriate, and intelligent fashion. Compromises will be essential, due to limits in computation, throughput of display hardware, and demands of real-time interaction, but our algorithms aim to balance the physical device constraints with carefully crafted models, general solutions, and thoughtful organization.

This study will tend to focus on one particularly well-motivated application for human models: human factors analysis. While not as exciting as motion picture characters, as personable as cartoons, or as skilled as Olympic athletes, there are justifiable uses to virtual human figures in this domain. Visualizing the appearance, capabilities and performance of humans is an important and demanding application (Plate 1). The lessons learned may be transferred to less critical and more entertaining uses of human-like models. From modeling realistic or at least reasonable body size and shape, through the control of the highly redundant body skeleton, to the simulation of plausible motions, human figures offer numerous computational problems and constraints. Building software for human factors applications serves a widespread, non-animator user population. In fact, it appears that such software has broader application since the features needed for analytic applications – such as multiple

simultaneous constraints – provide extremely useful features for the conventional animator. Our software design has tried to take into account a wide variety of physical problem-oriented tasks, rather than just offer a computer graphics and animation tool for the already skilled or computer-sophisticated animator.

The remainder of this chapter motivates the human factors environment and then traces some of the relevant history behind the simulation of human figures in this and other domains. It concludes with a discussion of the specific features a human modeling and animation system should have and why we have concentrated on some and not others. In particular, we are not considering cognitive problems such as perception or sensory interpretation, target tracking, object identification, or control feedback that might be important parts of some human factors analyses. Instead we concentrate on modeling a virtual human with reasonable biomechanical structure and form, as described in Chapter 2. In Chapter 4 we address the psychomotor behaviors manifested by such a figure and show how these behaviors may be interactively accessed and controlled. Chapter 5 presents several methods of motion control that bridge the gap between biomechanical capabilities and higher level tasks. Finally, in Chapter 6 we investigate the cognition requirements and strategies needed to have one of these computational agents follow natural language task instructions.

1.1 Why Make Human Figure Models?

Our research has focused on software to make the manipulation of a simulated human figure easy for a particular user population: human factors design engineers or ergonomics analysts. These people typically study, analyze, assess, and visualize human motor performance, fit, reach, view, and other physical tasks in a workplace environment. Traditionally, human factors engineers analyze the design of a prototype workplace by building a mock-up, using real subjects to perform sample tasks, and reporting observations about design satisfaction. This is limiting for several reasons. Jerry Duncan, a human factors engineer at Deere & Company, says that once a design has progressed to the stage at which there is sufficient information for a model builder to construct the mock-up, there is usually so much inertia to the design that radical changes are difficult to incorporate due to cost and time considerations. After a design goes into production, deficiencies are alleviated through specialized training, limits on physical characteristics of personnel, or various operator aids such as mirrors, markers, warning labels, etc. The goal of computer-simulated human factors analysis is not to replace the mock-up process altogether, but to incorporate the analysis into early design stages so that designers can eliminate a high proportion of fit and function problems before building the mock-ups. Considering human factors and other engineering and functional analyses together during rather than after the major design process is a hallmark of Concurrent Engineering [Hau89].

It is difficult to precisely characterize the types of problems a human factors engineer might address. Diverse situations demand empirical data on human capabilities and performance in generic as well as highly specific tasks. Here are some examples.

- Population studies can determine body sizes representative of some group, say NASA astronaut trainees, and this information can be used to determine if space vehicle work cells are adequately designed to fit the individuals expected to work there. Will all astronauts be able to fit through doors or hatches? How will changes in the workplace design affect the fit? Will there be unexpected obstructions to zero gravity locomotion? Where should foot- and hand-holds be located?
- An individual operating a vehicle such as a tractor will need to see the surrounding space to execute the task, avoid any obstructions, and insure safety of nearby people. What can the operator see from a particular vantage point? Can he control the vehicle while looking out the rear window? Can he see the blade in order to follow an excavation line?
- Specific lifting studies might be performed to determine back strain limits for a typical worker population. Is there room to perform a lift properly? What joints are receiving the most strain? Is there a better posture to minimize torques? How does placement of the weight and target affect performance? Is the worker going to suffer fatigue after a few iterations?
- Even more specialized experiments may be undertaken to evaluate the comfort and feel of a particular tool's hand grip. Is there sufficient room for a large hand? Is the grip too large for a small hand? Are all the controls reachable during the grip?

The answers to these and other questions will either verify that the design is adequate or point to possible changes and improvements early in the design process. But once again, the diversity of human body sizes coupled with the multiplier of human action and interaction with a myriad things in the environment leads to an explosion in possible situations, data, and tests.

Any desire to build a "complete" model of human behavior, even for the human factors domain, is surely a futile effort. The field is too broad, the literature immense, and the theory largely empirical. There appear to be two directions out of this dilemma. The first would be the construction of a computational database of all the known, or at least useful, data. Various efforts have been undertaken to assemble such material, for example, the NASA sourcebooks [NAS78, NAS87] and the *Engineering Data Compendium* [BKT86, BL88]. The other way is to build a sophisticated computational human model and use it as a subject in simulated virtual environment tests. The model will utilize an ever-expanding human factors data set to dictate its performance. Upon some reflection, it appears that database direction may