

Artificial Animals for Computer Animation:

Biomechanics, Locomotion, Perception, and Behavior

by

涂晓媛

Xiaoyuan Tu



A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Computer Science
University of Toronto

© Copyright by Xiaoyuan Tu 1996

© Xiaoyuan Tu 1996
ALL RIGHTS RESERVED

Artificial Animals for Computer Animation: Biomechanics, Locomotion, Perception, and Behavior

Doctor of Philosophy, 1996

Xiaoyuan Tu

Department of Computer Science, University of Toronto

Abstract

This thesis develops an artificial life paradigm for computer graphics animation. Animals in their natural habitats have presented a long-standing and difficult challenge to animators. We propose a framework for achieving the intricacy of animal motion and behavior evident in certain natural ecosystems, with minimal animator intervention.

Our approach is to construct artificial animals. We create self-animating, autonomous agents which emulate the realistic appearance, movement, and behavior of individual animals, as well as the patterns of social behavior evident in groups of animals. Our computational models achieve this by capturing the essential characteristics common to all biological creatures—biomechanics, locomotion, perception, and behavior.

To validate our framework, we have implemented a virtual marine world inhabited by a variety of realistic artificial fishes. Each artificial fish is a functional autonomous agent. It has a physics-based, deformable body actuated by internal muscles, sensors such as eyes, and a brain with perception, motor, and behavior control centers. It swims hydrodynamically in simulated water through the controlled coordination of its muscle actions. Artificial fishes exhibit a repertoire of behaviors. They explore their habitat in search of food, navigate around obstacles, contend with predators, and indulge in courtship rituals to secure mates. Like their natural counterparts, their behavior is based on their perception of the dynamic environment and their internal motivations.

Since the behavior of artificial fishes is adaptive to their virtual habitat, their detailed motions need not be keyframed nor scripted. This thesis therefore demonstrates a powerful approach to computer animation in which the animator plays the role of a nature cinematographer, rather than the more conventional role of a graphical model puppeteer. Our work not only achieves behavioral animation of unprecedented complexity, but it also provides an interesting experimental domain for related research disciplines in which functional, artificial animals can serve as autonomous virtual robots.

Acknowledgment

First, I would like to thank my advisors Demetri Terzopoulos and Eugene Fiume. This thesis would not have been possible without their support.

Demetri has been my mentor in the exciting fields of computer animation and physics-based graphics modeling. His many remarkable qualities have benefited me greatly. I am most impressed by his breadth of knowledge, his ability to seize new ideas, his openmindedness, and his perfectionism towards technical writing. I am especially grateful for the unfailing encouragement he has given me through the years. I would not have come this far this quickly without his expert guidance and the substantial time and effort he has put into my thesis.

I thank Eugene for giving me the chance to join the graphics group three years ago. Without his kindness, I would have been stuck studying boring operating systems (no offense intended). I appreciate his many insightful comments that have helped to improve my thesis and to put my view into perspective.

My special thanks go to Professor Daniel Thalmann for generously agreeing to serve as external examiner of my thesis, despite his busy schedule and the long distance he had to travel. Special thanks also go to Professor Janet Halperin for serving as internal examiner and Professor Michiel Van de Panne for serving as internal appraiser on my thesis committee. Many thanks to Professors Ken Sevcik and Geoffrey Hinton for serving on my committee and providing valuable comments on my thesis.

I thank the wonderful people of the graphics and vision labs. My years at U of T would have been dull without them. I would like to express my sincere gratitude to Tamara Stephas, Meng Sun, Victor Ng, Xiaohuan Wang, Petros Faloutsos, Jeremy Cooperstock, Radek Grzeszczuk, Beverly Harrison, Kevin Schlueter, Michael McCool, Alex Mitchell, Jeff Tupper, Jos Stam, Sherif Ghali, Joe Laszlo, William Hunt and Baining Guo from the graphics lab, and Tim McInerney, Victor Lee and Xuan Ju from the vision lab. My deep appreciation to my friends from other groups in the department, Vincent Gogan, Steven Shapiro and Michalis Faloutsos, and, from elsewhere, Yuxiang Wang, Wei Xu, Tina Shapiro and Christopher Lori. Particularly, I thank my colleague Radek Grzeszczuk, for

collaborating in the production of “Cousto”. Thanks also to Kathy Yen for helping me to deal with the many tedious details in arranging my defense.

I would especially like to thank Michiel Van de Panne, who was a student at the graphics lab when I started my PhD program and now, as I finish, is a professor. He has always been very generous in helping others around the lab. He taught me how to use the video equipment and provided substantial assistance in the production of my animations.

I thank my parents, to whom I dedicate this thesis, for their endless love and support, for their unwavering confidence in me and their limitless patience. They are the best parents I could have ever dreamed to have.

Finally, thanks to my partner John, for his love, support, help (both as a colleague and a friend) and the wonderful times he brings me.

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 1 |
| 1.1 | Motivation | 1 |
| 1.2 | Challenges | 2 |
| 1.2.1 | Conventional Animation Techniques | 3 |
| 1.3 | Methodology: Artificial Life for Computer Animation | 5 |
| 1.3.1 | Criteria and Goals | 5 |
| 1.3.2 | Artificial Animals | 5 |
| 1.3.3 | From Physics to Realistic Locomotion | 6 |
| 1.3.4 | Realistic Perception | 8 |
| 1.3.5 | Realistic Behavior | 8 |
| 1.3.6 | Fidelity and Efficiency | 9 |
| 1.4 | Contributions and Results | 10 |
| 1.4.1 | Primary Contributions | 11 |
| 1.4.2 | Auxiliary Technical Contributions | 14 |
| 1.5 | Thesis Overview | 15 |
| 2 | Background | 17 |
| 2.1 | Physics-Based Modeling | 17 |
| 2.1.1 | Constraint-based Approach | 18 |
| 2.1.2 | Motion Synthesis Approach | 19 |
| 2.2 | Behavioral Animation | 22 |

| | | |
|----------|---|-----------|
| 2.2.1 | Perception Modeling | 23 |
| 2.2.2 | Control of Behavior | 24 |
| 2.3 | The Modeling of Action Selection | 25 |
| 2.3.1 | Defining Action | 25 |
| 2.3.2 | Goals and Means | 27 |
| 2.3.3 | Previous Work | 28 |
| 2.3.4 | Task-level Motion Planning | 30 |
| 2.4 | Summary | 30 |
| 3 | Functional Anatomy of an Artificial Fish | 33 |
| 3.1 | Motor System | 35 |
| 3.2 | Perception System | 35 |
| 3.3 | Behavior System | 35 |
| 4 | Biomechanical Fish Model and Locomotion | 37 |
| 4.1 | Discrete Physics-Based Models | 37 |
| 4.2 | Structure of the Dynamic Fish Model | 38 |
| 4.3 | Mechanics | 41 |
| 4.3.1 | Viscoelastic Units | 41 |
| 4.4 | Muscles and Hydrodynamics | 44 |
| 4.5 | Numerical Solution | 46 |
| 4.5.1 | System Matrix Assembling and the Skyline Storage Scheme | 48 |
| 4.5.2 | Algorithm Outline and Discussion | 51 |
| 4.6 | Motor Controllers | 52 |
| 4.6.1 | Muscle Motor Controllers | 53 |
| 4.6.2 | Pectoral Fin Motor Controllers | 57 |
| 5 | Modeling the Form and Appearance of Fishes | 59 |
| 5.1 | Constructing 3D Geometric Fish Models | 59 |

| | | |
|----------|---|-----------|
| 5.2 | Obtaining Texture Coordinates | 62 |
| 5.2.1 | Deformable Mesh | 63 |
| 5.3 | Texture-Mapped Models | 64 |
| 5.4 | Coupling the Dynamic and Display Models | 65 |
| 5.5 | Visualization of the Pectoral Motions | 67 |
| 6 | Perception Modeling | 71 |
| 6.1 | Perception Modeling for Animation | 71 |
| 6.2 | Overview of the Artificial Fish's Perception System | 73 |
| 6.3 | Vision Sensor Modeling | 73 |
| 6.3.1 | Perceptual Range | 74 |
| 6.3.2 | Occlusion | 75 |
| 6.3.3 | Functionality | 75 |
| 6.4 | Computing Visibility | 77 |
| 6.4.1 | Visibility of a Point | 77 |
| 6.4.2 | Visibility of Another Fish | 80 |
| 6.4.3 | Visibility of a Cylinder | 80 |
| 6.4.4 | Visibility of Seaweeds | 80 |
| 6.4.5 | Discussion | 81 |
| 6.5 | The Focuser | 82 |
| 6.5.1 | Focus of Attention in Animals | 82 |
| 6.5.2 | Design of the Focuser | 83 |
| 6.5.3 | Summary | 86 |
| 6.6 | From Perception to Behavior | 87 |
| 6.6.1 | An Example: Collision Detection | 87 |
| 6.7 | Synthetic Vision Models | 90 |
| 7 | The Behavior System | 93 |

| | | |
|----------|--|------------|
| 7.1 | Effective Action Selection Mechanisms | 94 |
| 7.2 | Behavior Control and Ethology | 95 |
| 7.2.1 | The Intention Level | 95 |
| 7.2.2 | The Action Level | 95 |
| 7.2.3 | Summary | 97 |
| 7.3 | Habits | 98 |
| 7.4 | Mental State | 98 |
| 7.5 | Intention Generator | 101 |
| 7.5.1 | Why Hierarchy? | 103 |
| 7.6 | Intention-Guided Perception: Control of the Focusser | 104 |
| 7.7 | Persistence in Behavior | 105 |
| 7.7.1 | Behavior Memory | 105 |
| 7.7.2 | Inhibitory Gain and Fatigue | 106 |
| 7.7.3 | Persistence in Targeting | 106 |
| 7.8 | Behavior Routines | 107 |
| 7.8.1 | Primitive Behavior: Avoiding Potential Collisions | 108 |
| 7.8.2 | Primitive Behavior: Moving Target Pursuit | 108 |
| 7.9 | Artificial Fish Types | 112 |
| 7.9.1 | Predators | 112 |
| 7.9.2 | Prey | 114 |
| 7.9.3 | Pacifists | 119 |
| 7.10 | Discussion | 124 |
| 7.10.1 | Analysis | 124 |
| 7.10.2 | Summary | 126 |
| 8 | Modeling the Marine Environment | 129 |
| 8.1 | Water Current | 129 |
| 8.2 | Seaweeds and Plankton | 130 |

| | |
|---|------------|
| 9 The Graphical User Interface | 133 |
| 9.1 Initialization Panels | 133 |
| 9.2 Manipulation Panels | 134 |
| 9.3 Control Panels | 137 |
| 9.4 Discussion | 138 |
| 10 Animation Results | 141 |
| 10.1 “Go Fish!” | 141 |
| 10.2 “The Undersea World of Jack Cousto” | 144 |
| 10.3 Animation Short: Preying Behavior | 144 |
| 10.4 Animation Short: Schooling Behavior | 146 |
| 10.5 Animation Short: Mating Behavior | 147 |
| 11 Conclusion and Future Work | 149 |
| 11.1 Conclusion | 149 |
| 11.2 Additional Impact in Animation and Artificial Life | 150 |
| 11.3 Impact in Computer Vision and Robotics | 152 |
| 11.4 Potential Applications in Ethology | 153 |
| 11.5 Other Artificial Animals | 154 |
| 11.6 Future Research Directions | 155 |
| 11.6.1 Animation | 155 |
| 11.6.2 Artificial Life | 157 |
| A Deformable Contour Models | 158 |
| B Visualization of the Pectoral Motions | 161 |
| B.1 Animating the Pectoral Flapping Motion | 161 |
| B.2 Animating the Pectoral Oaring Motion | 163 |
| C Prior Action Selection Mechanisms | 164 |

| | |
|---------------------------------------|------------|
| C.1 Behavior Choice Network | 164 |
| C.2 Free-Flow Hierarchy | 165 |
| D Color Images | 173 |

List of Figures

| | | |
|------|---|----|
| 1.1 | Artificial fishes in their physics-based world. See the original color image in Appendix D . | 12 |
| 1.2 | Mating behavior. Female (top) is courted by larger male. See the original color image in Appendix D | 12 |
| 1.3 | Predator shark stalking school of prey fish. See the original color image in Appendix D . | 13 |
| 2.1 | Differentiating action selection from motor control in design. | 26 |
| 3.1 | System overview of the artificial fish. | 34 |
| 4.1 | The motor system of an artificial fish. | 38 |
| 4.2 | The biomechanical fish model. | 40 |
| 4.3 | The top and side view of the outline of the fish model at rest. | 40 |
| 4.4 | Uniaxial elastic, viscous, and viscoelastic units. | 42 |
| 4.5 | Hydrodynamic locomotion. | 45 |
| 4.6 | Structure of the system matrix. | 49 |
| 4.7 | The skyline storage scheme. | 50 |
| 4.8 | Top view of an artificial fish during caudal swimming motion. | 55 |
| 4.9 | Photograph of real fish swimming. | 55 |
| 4.10 | Top-front view of an artificial fish during caudal swimming motion. | 55 |
| 4.11 | The steering map. | 56 |
| 4.12 | The turning motion of the artificial fish. | 56 |
| 4.13 | The pectoral fin geometry. | 57 |
| 4.14 | The local coordinate system of a fish. | 58 |

| | | |
|------|---|-----|
| 5.1 | Digital images of real fish. | 60 |
| 5.2 | Control-point mesh of the left half of a fish body. | 61 |
| 5.3 | The shaded NURBS surface of the left half of a fish body. | 61 |
| 5.4 | Geometric model of fish body with fins. | 62 |
| 5.5 | Control-point mesh fish models. | 62 |
| 5.6 | Initial and stabilized deformable mesh. | 64 |
| 5.7 | Texture mapped 3D fish models. | 64 |
| 5.8 | The subdivision of the faces of the dynamic fish. | 66 |
| 5.9 | Geometric mesh surface overlayed on dynamic fish model. | 68 |
| 5.10 | The geometric NURBS surface fish deforms with the dynamic fish. | 68 |
| 5.11 | Snapshots of the pectoral flapping motion. | 69 |
| 5.12 | Snapshots of the pectoral oaring motion of a butterfly fish. | 69 |
| | | |
| 6.1 | The perception system in an artificial fish. | 73 |
| 6.2 | Vision sensor. | 74 |
| 6.3 | Occlusion and distance limit visual perception. | 75 |
| 6.4 | Fisheye view of the world showing fishing line. See the original color image in Appendix D. | 76 |
| 6.5 | Fisheye view of the world showing hooked fish. | 77 |
| 6.6 | Occlusion test. | 78 |
| 6.7 | Swimming fishes avoiding collisions with cylindrical obstacles. | 88 |
| 6.8 | The fish's collision sensitivity region. | 90 |
| 6.9 | Analogues of intrinsic images. | 92 |
| | | |
| 7.1 | The behavior system in an artificial fish. | 94 |
| 7.2 | The control scheme for action selection. | 96 |
| 7.3 | The form of the stimulus functions $S^h(t)$ and $S^l(t)$ | 100 |
| 7.4 | Generic intention generator. | 102 |
| 7.5 | Outline of <i>chasing-target</i> routine. | 109 |

| | | |
|------|--|-----|
| 7.6 | The position of target relative to fish. | 109 |
| 7.7 | A peaceful marine world. See the original color image in Appendix D. | 111 |
| 7.8 | The smell of danger. See the original color image in Appendix D. | 111 |
| 7.9 | The intention generator of a predator. | 112 |
| 7.10 | The angle E_k | 113 |
| 7.11 | A hungry predator ingesting prey. | 114 |
| 7.12 | Portion of intention generator of prey. | 114 |
| 7.13 | Schooling behavior routine. | 115 |
| 7.14 | A small school of angelfish. | 116 |
| 7.15 | An inceptive school. | 116 |
| 7.16 | A school of fish encounters an obstacle. | 117 |
| 7.17 | School scatters in terror. | 118 |
| 7.18 | Fleeing from predator. | 118 |
| 7.19 | Mating behaviors in real fishes. | 121 |
| 7.20 | Two males (larger) vying for the attention of a female. | 122 |
| 7.21 | A clownfish displays opportunism in behavior. | 126 |
| 8.1 | The flow primitives. | 129 |
| 8.2 | The swaying motion of a seaweed in simulated water currents. | 131 |
| 8.3 | The seaweed responds to hydrodynamic forces produced by a passing fish. | 132 |
| 9.1 | The object initialization panel. | 134 |
| 9.2 | The position initialization panel. | 135 |
| 9.3 | The physical parameter panel. | 135 |
| 9.4 | The behavior panel. | 136 |
| 9.5 | The habit panel. | 137 |
| 9.6 | The general control panel. | 138 |
| 9.7 | The fishview control panel and the rendered identity maps from fish '1'. | 139 |

| | | |
|------|---|-----|
| 10.1 | Denizens of the virtual marine world happily feeding on plankton. | 142 |
| 10.2 | Hungry fishes approaching the hook. | 143 |
| 10.3 | A hapless fish is caught and pulled to the surface. | 143 |
| 10.4 | A dangerous shark stalking a school of prey. | 145 |
| 10.5 | A school of fleeing prey. | 145 |
| 10.6 | A pair of courting pokerfish. | 146 |
| 10.7 | A large predator hunting small prey fishes while pacifist fishes, untroubled by the predator, feed on plankton. | 147 |
| 11.1 | Locomotion learning center in the brain of the artificial fish. | 151 |
| B.1 | Visualization of pectoral fin motion. | 162 |

Chapter 1

Introduction

A kangaroo hops across a barren plain. It leaps and lands rhythmically. The harmonic movements of its body, legs and tail trace out graceful curves in the air. A flock of birds glides across the sky. Individual birds flap their wings and adjust their direction autonomously, yet they all fly in unison.

1.1 Motivation

Animals in motion have intrigued computer graphics animators and researchers for several decades. They have long been the subject of study of zoologists and ethologists, and have recently helped inspire the emerging research discipline of artificial life.

In computer graphics, most animations of animals have been created using the traditional and often highly labour intensive keyframing technique, in which computers are employed to interpolate between animator-specified keyframes (Lasseter, 1987). More recently, increasingly automated techniques for synthesizing realistic animal motion have drawn much attention. Successful attempts have been made to animate the motions of humans (Magenat-Thalmann and Thalmann, 1990; Hodgins, Sweeney and Lawrence, 1992), of certain animals (Miller, 1988; Girard, 1991) and of some imaginary creatures (Witkin and Kass, 1988; van de Panne and Fiume, 1993; Ngo and Marks, 1993). However, motion synthesis is only part of the challenge of animating animals. Some group behaviors evident in the animal world, such as flocking, schooling and herding (Reynolds, 1987) have also been simulated and realistically animated in recent feature films.

In this dissertation, we will investigate the problem of producing animation which captures

the intricacy of motion and behavior evident in certain natural ecosystems. These animations are intrinsically complex and present a challenge to the computer graphics practitioner. Animations of this sort are of interest not only because they attempt to recreate fascinating natural scenarios, but also because they have broad applicability. They can be used in the entertainment industry, for special effects in movies, for video games, for virtual reality rides; as well as in education as, say, interactive educational tools for teaching biology.

Our goal will be to create the animations that we have described, not by conventional keyframing, but rather through the sophisticated modeling of animals and their habitats. To this end, we have been motivated by and have contributed to the artificial life (ALife) movement (Levy, 1992). ALife complements the traditional analytic approach of biology by aiming to understand natural life through synthetic, computational means. That is to say, rather than studying biological phenomena by analyzing living systems, the ALife approach attempts to synthesize artificial systems that behave like living organisms. An important area of ALife research is the synthesis of artificial animals—or “animats”—implemented both in software and in hardware (Meyer and Guillot, 1991; Cliff et al., 1994). Computational models of simple animals, such as single-cell life forms (Langton, 1987) and insects (Beer, 1990), have been proposed with interesting results. Many of these models draw upon theories of animal behavior put forward by ethologists (Manning, 1979; McFarland, 1971).

Since we will view the animation of natural ecosystems as the process of visualizing computer simulations of animals in their habitats, our work straddles the boundary between the fields of computer graphics and artificial life. This theme has also been investigated by Terzopoulos *et al.* (1995).

1.2 Challenges

Natural ecosystems are as challenging to animate as they are fascinating to watch. The major challenge comes from their intrinsic complexity. In a given animation system, there may be a large number of animals, each of which may exhibit elaborate behaviors. Ideally, one would like to achieve an abundance of natural intricacy with minimal effort on the part of the animator. The challenge is exacerbated when one also demands visual authenticity in the appearance and locomotion of individual animals and in their behavior.

Ecosystems are characterized by the relationships between animals and their habitats. This means that when we look at an ecosystem, we are keenly aware of the behaviors exhibited by animals as

they interact with their dynamic environment, especially with other animals. For example, a hunting scenario will not seem authentic if a rabbit hops around carelessly disregarding the presence of a hungry wolf, or if a crowd of pigeons rest calmly while a child runs in their midst. People's familiarity with various animals imposes strict criteria for the evaluation of the visual results of our proposed simulations, since even small imperfections in the animated motions or behaviors will be readily recognizable.

Visual realism, however, is not the only constraint on such animations. For applications in the entertainment and educational industries, the animator should be able to control various aspects of an animation. It is especially important to be able to easily modify an animation; for example, altering the virtual environment, changing the number, types and distribution of the virtual animals and, moreover, varying the personalities of the animals and even interacting with them.

1.2.1 Conventional Animation Techniques

Traditional computer animation techniques, such as keyframing, have been used to create many great animations, including those of animals. However, they have several limitations:

Significant Animator Intervention is Required

Perhaps the most spectacular instance to date of conventional animation techniques applied to the animation of animals is the dinosaurs in the blockbuster feature film "*Jurassic Park*" (a 1993 Amblin Entertainment Production for Universal Pictures). Yet as realistic looking as they may be, these dinosaurs are merely graphical puppets which require teams of highly skilled human animators to plot their actions and detailed motions carefully from one step to the next. This reveals the main drawback of keyframing: The amount of effort expended by the animator increases dramatically with the length, complexity and intended realism of the animation.

Techniques that do not require as much animator skill, such as motion capture schemes (Calvert, Chapman and Patla, 1980), have also been widely used in producing realistic animated motions. However, they tend to be inflexible, since they produce motions that are highly specific, hard to parameterize, and difficult to compose into lengthier animations. Moreover, such schemes are not easily applied to non-human or imaginary creatures.

Characters Lack Autonomy

The lack of autonomy of keyframed or motion-data driven characters detracts from the modifiability and interactability of the animation. Since realistic behaviors of an animal are subject to the state of its environment (which contains immobile and mobile objects, such as trees, stones, and other animals), slight modifications in a scripted animation, such as moving a tree or adding another animal, may require that the entire animation be re-scripted. Graphical puppets are incapable of actively interacting with their environment. As a result, keyframing techniques are ill-suited to applications such as virtual reality, computer games and interactive educational tools.

Low-Level Motion Specification is Burdensome

Using conventional animation techniques, animations are specified at a very low level, thus granting the animator complete control over every aspect of the animation. Complete animator control is sometimes required, especially in cartooning; however, it is generally unnecessary for the sorts of applications in which we are interested. Consider the case of animating a realistic virtual zoo: It is not important if some particular animal appears in some specific posture at some specific time instant; what is important is for the lions and monkeys to look real and for them to move and behave realistically. Complete animator control is also undesirable in our case because it implies little or no autonomy in the animated characters themselves. Therefore, a good strategy for our purposes is to relinquish low-level animator control in favor of a much higher level of control.

Physical Realism is Not Guaranteed

Last but not least, using conventional keyframing techniques, realism depends solely upon the skills of the animator. Unless the animator is highly skilled, poor visual results may obtain. In particular, since conventional geometric models possess no mechanical properties, the physical correctness of the resulting motion is not guaranteed, nor will the animated figures respond to forces in a realistic manner.

1.3 Methodology: Artificial Life for Computer Animation

1.3.1 Criteria and Goals

In light of the preceding discussion, we seek an approach to animating natural ecosystems that is capable of achieving realistic visual effects through an automatic process. The desired properties of such an approach are as follows:

1. The appearance, locomotion and behavior of the animated creatures should be visually convincing.
2. The creatures should have a high degree of autonomy so that this can be achieved with minimal animator intervention. The level of autonomy in the animated animals should, however, permit and support the necessary degree of high-level animator control:
 - The animator should be able to alter the initial conditions of the animation, such as the number and positions of immobile and mobile objects in the virtual habitat.
 - The animator should be able to influence or direct the behaviors of the animated characters to some degree.

The research reported in this thesis develops a highly automatic approach to creating life-like animations and validates it through implementation. Realism is achieved through advanced modeling of animals and their habitats.

1.3.2 Artificial Animals

We believe that the best way of achieving our goals in the long run is to pursue the challenging approach of constructing *artificial animals*. The properties and internal control mechanisms of artificial animals should be qualitatively similar to those of their natural counterparts.

There are several properties common to all animals. The most salient one is that *all animals are autonomous*: They have physical bodies that are actuated by muscles, enabling them to locomote; they have eyes and other sensors to actively perceive their environment; they have brains which interpret their perceptions and govern their actions. Indeed, autonomy is the consequence of possessing a brain capable of controlling perception and action in a physical body. The behavior of

an animal is a consequence of its autonomous interaction with its environment to satisfy its survival needs. No external control is required for animals to cope with their dynamic habitats, yet their autonomy does not prevent higher animals from being influenced or directed (consider trained circus animals, or human actors).

Artificial animals should be self-animating actors that emulate the autonomy of real animals. In our artificial life approach to computer animation, we build animal-like autonomy into our graphics models; not only to minimize the amount of animator intervention while supporting modifiability and interactability, but also to obtain behavioral realism. Our main concern is how to model the locomotion, perception and behavior capabilities of animals and how to integrate these models effectively within a life-like artificial animal. Our research in this respect shares common goals with ALife research, where artificial animals are often referred to as “animats”. Previous animats have been models of simple creatures and the behaviors simulated usually pertain to genetic reproduction and “natural selection” (Langton, 1987; Varela and Bourgine, 1991). We attempt to develop artificial life patterned after animals that are more evolved and have a significantly broader range of behavior.

In the following sections, we will identify the essential properties and mechanisms that allow real animals to locomote effectively, to perceive, and hence to behave autonomously. From this we derive design methodologies for achieving realistic animated locomotion and behavior with minimal animator intervention.

1.3.3 From Physics to Realistic Locomotion

Physics-Based Modeling

The motion of any physical entity is governed at the lowest level by the laws of physics. The use of physics is not new to computer graphics. It was introduced as “physics-based modeling” about a decade ago (Armstrong and Green, 1985; Wilhelms, 1987; Terzopoulos et al., 1987) and has spawned a large body of advanced graphics modeling and animation research. Using physics-based models for graphics not only ensures physical realism of the resulting motion, it also allows subtle yet visually important motions to be animated automatically. Consider, for example, the realistic animation of a stampeding elephant. It would be a heroic chore to try to apply manual keyframing or other purely kinematic methods to animate the rippling flesh, the flapping ears, or the swinging trunk and tail. Physics-based modeling is capable of producing such motions automatically. More details about physics-based modeling and related previous work can be found in Chapter 2.

Simulated Physical Body

The laws of physics and the principles of biomechanics shape the appearance of animal motion. Therefore, the best way to achieve realistic locomotion is to simulate their effects. According to mechanics, the change of the state of an object, or what we commonly call “movement”, is caused by forces. The motion of an inanimate object, such as a stone, is generally caused by unbalanced external forces and hence is *passive*. The motion of an animal, on the other hand, is generally initiated by unbalanced internal forces actively generated by its muscles and hence is *active*. Each species of animal has its particular body structure and arrangement of muscles. This in turn dictates its particular mode of locomotion. We therefore construct simulated physical bodies for our artificial animals. By doing so, the laws of physics will guarantee the physical correctness of the resulting motion, while the biomechanical principles relevant to the animal of interest can yield natural muscle actions that simulate the particular locomotion patterns characteristic of the animal in its physical environment.

Simulated Physical Environment

The locomotion of an animal not only derives from the dynamics of its body but also is a result of the dynamics of its environment. In accordance with biomechanical principles, various locomotion patterns of animals—e.g. flying, swimming or running—emerge from the interaction between their active muscle-actuated body movements and the reactive physical environment—air, water, or *terra firma*. For example, birds flap their wings inducing aerodynamic forces, which in turn enable them to fly through the air. They cannot fly in a vacuum. Therefore, in addition to modeling the physics of animal bodies, we need also to model the physics of their environments.

Motor Control

Upon the computational physics substrate, computed via a dynamic simulation, it is possible to generate realistic locomotion of the artificial animal through simulated muscle control. Through evolution, most natural animals have developed their particular mechanisms for motor control, where coordinated muscle activations result in energy-efficient locomotion. Effective motor controllers can be derived based on the *a priori* knowledge about the characteristic muscle activations of the corresponding real animal.

1.3.4 Realistic Perception

In order to survive in dynamic and often hostile environments, animals are able to adapt their behaviors according to the current situation. As summarized by the ethologist Manning (1979), the behavior of an animal “includes all those processes by which the animal senses the external world and the internal state of its body and responds to changes which it perceives.” This definition emphasizes the crucial dependence of animal behavior on perception, for without perception, an animal cannot possibly react to its environment. To increase their chances of survival, most animals have evolved acute perceptual modalities, especially eyes, to detect opportunities and dangers in their habitat. The sense organs of animals have specific capabilities and inherent limitations. For example vision is most effective within some proximal distance because, under most circumstances, spatially proximal events will have the greatest effect on an animal. Furthermore, vision is not possible through opaque objects. We must model both the capabilities and the limitations of perception correctly for our artificial animals to exhibit realistic behavior.

1.3.5 Realistic Behavior

Given that an animal has the ability to locomote and to sense its environment, its brain is able to interpret the sensory information and select appropriate actions to yield a useful range of behavior.

Environment, External Stimuli and Internal State

To enable an artificial animal to behave realistically and autonomously, it is necessary to model relevant aspects of its habitat as well as its internal mental state. Sensory stimuli present information about environmental events such as the presence of food, which may cause the animal to ingest, or the presence of a predator, which may cause the animal to flee. However, external stimuli alone cannot fully determine an animal’s behavior. An animal that is satiated will normally not ingest more food even if food is available. If an animal is desperately thirsty, it may delay taking evasive action despite the presence of a predator in the distance in order to drink at a waterhole. Its decision to engage in a particular behavior is predicated on the internal state of the animal which reflects the physical condition of its body—whether it is hungry, tired, etc. Such internal state can thus be considered as inducing the “need” or “motivation” to evoke a specific behavior.

Action Selection

After an animal obtains sensory information about its external world and internal state, it has to process this information to decide what to do next. In particular, the perceived external stimuli must be evaluated with respect to the animal's internal state in order for it to determine the most appropriate course of action. This higher control process is often referred to as "action selection." The brain can carry out the action selection process at the cognitive or sub-cognitive level. The action selection mechanism is the key to adaptiveness and autonomy. It is essential to design effective action selection mechanisms for artificial animals.

Behavioral Animation

The modeling methodology that we have outlined in the previous sections may be viewed as a sophisticated form of *behavioral animation*, in which autonomous models are built by introducing perception and certain behavioral components into the motion control algorithms (Reynolds, 1987). Interestingly, during the last decade, much of the attention in the graphics community has centered on realistic low-level motion synthesis, with only a few researchers pursuing the modeling of realistic behavior. Prior behavioral animation work, however, paid little attention to the realism of the motion of individual creatures. Also, prior work was generally restricted to the animation of one or two specific behaviors and not to the development of broad behavioral repertoires.

1.3.6 Fidelity and Efficiency

Our methodology aims at producing autonomous artificial animals that not only look like, but also move like and behave like their natural counterparts. An important question is: How closely should our models attempt to emulate real animals? Clearly, a certain level of modeling fidelity is required in order to generate convincing results and, generally, the more faithful the models, the more realistic the results. Most real animals of interest are extremely complex both in terms of body structure and in terms of behavioral repertoire. Models of animals can therefore easily become excessively complicated. However, it is desirable for an animation system to be reasonably efficient so that it will run quickly enough on current graphics computers to allow interactive modification by the animator.

For the purposes of animation, we must strike a good compromise between model fidelity and

computational efficiency. Striking the proper balance is a critical design issue, since inappropriate model accuracy can be counterproductive to the purpose at hand. For example, if we wanted to build a model of a tiger to show the effect of gait on the maturation of the bone in its legs, it may be necessary to model the cellular structure of the bone. However, this is hardly necessary if we are only interested in animating tiger gaits. Therefore we should keep the model complexity as low as is necessary to achieve the intended purpose—in our case, realistic appearance, locomotion and behavior.

1.4 Contributions and Results

Fishes¹, the superclass Pisces, are an important species of animals that exhibit elaborate behaviors both as individuals and in groups. Most people find them fascinating to watch. Their behavioral complexity is generally higher than that of most insects, but lower than that of most mammals. Since realistic, automatic animation of the locomotion of humans and other advanced animals remains elusive, a visually convincing virtual marine world presents an excellent choice for validating our approach to animating natural ecosystems.

Imagine a virtual marine world inhabited by a variety of realistic fishes. In the presence of underwater currents, the fishes employ their muscles and fins to swim gracefully around static obstacles and among moving aquatic plants and other fishes. They autonomously explore their dynamic world in search of food. Large, hungry predator fishes stalk smaller prey fishes in the deceptively peaceful habitat. Prey fishes swim around contentedly until the sight of predators compels them to take evasive action. When a dangerous predator appears in the distance, similar species of prey form schools to improve their chances of survival. As the predator nears a school, the fishes scatter in terror. A chase ensues in which the predator selects victims and consumes them until satiated. Some species of fishes seem untroubled by predators. They find comfortable niches and feed on floating plankton when they get hungry. Driven by healthy libidos, they perform elaborate courtship rituals to secure mates.

We have successfully applied the basic methodology outlined in the previous section to develop an animation framework within whose scope fall all of the above complex patterns of behavior, and many more, without any keyframing. We aim to make the colorfully textured denizens of the

¹“Fish” is both singular and plural; when plural, it refers to more than one fish *within* the same species. The plural “fishes” is used when two or more species are involved (Wilson and Wilson, 1985).

fish world as realistic as the “Jurassic Park” dinosaurs. Yet, unlike the dinosaurs, each fish in this community exists as an independent, self-governing virtual agent. None of the actions is keyframed or scripted in advance, but is instead driven by the individual perceptions and internal desires of the artificial fishes. Each fish attends to a hierarchy of needs, with its brain considering the urgency of each situation. When the animation program is initiated, the operator specifies only which fish are present and their initial conditions. Upon starting the artificial life simulation, the creatures proceed to act of their own accord.

The visual results of this work are illustrated by two animations. Our 1993 animation “Go Fish!” (Tu, Terzopoulos and Fiume, 1993) shows a colorful variety of artificial fishes foraging in translucent water. A sharp hook on a line descends towards the hungry fishes and attracts them. A hapless fish, the first to bite the bait, is caught and drawn to the surface. The color plates show stills from our 1994 animation “The Undersea World of Jack Cousto” (Tu, Grzeszczuk and Terzopoulos, 1995). Fig. 1.1 shows a variety of animated artificial fishes. The reddish fish are engaged in a mating ritual, the large, dark colored fish is a predator hunting for small prey, the remaining fishes are feeding on plankton (white dots). Dynamic seaweeds grow from the ocean bed and sway in the current. In Fig. 1.2, the large male in the foreground is courtship dancing with the female (top). The prey fish in the background are engaging in schooling behavior, a common subterfuge for avoiding predators. Fig. 1.3 shows a shark stalking the school. The detailed motions of the artificial fishes emulate the complexity and unpredictability of movement of their natural counterparts, and this enhances the visual beauty of the animations.

1.4.1 Primary Contributions

This thesis contributes both to the field of computer graphics and to the field of artificial life. It leverages the synergy between these two fields for the realistic modeling, simulation, and animation of animals. Our contributions have been published in the computer graphics literature (Tu and Terzopoulos, 1994a; Tu and Terzopoulos, 1994b) and in the artificial life literature (Terzopoulos, Tu and Grzeszczuk, 1994a; Terzopoulos, Tu and Grzeszczuk, 1994b).

In computer graphics, our work is the first to combine within a unified framework extensive physics-based graphics models, locomotion control, and higher-level behavioral models for animation. In the context of artificial life, we develop animats of unprecedented realism and sophistication. Our life-like animations of fish in their habitat demonstrate a functional model that captures the interplay of physics, locomotion, perception and behavior in animals. The behaviors that our artificial

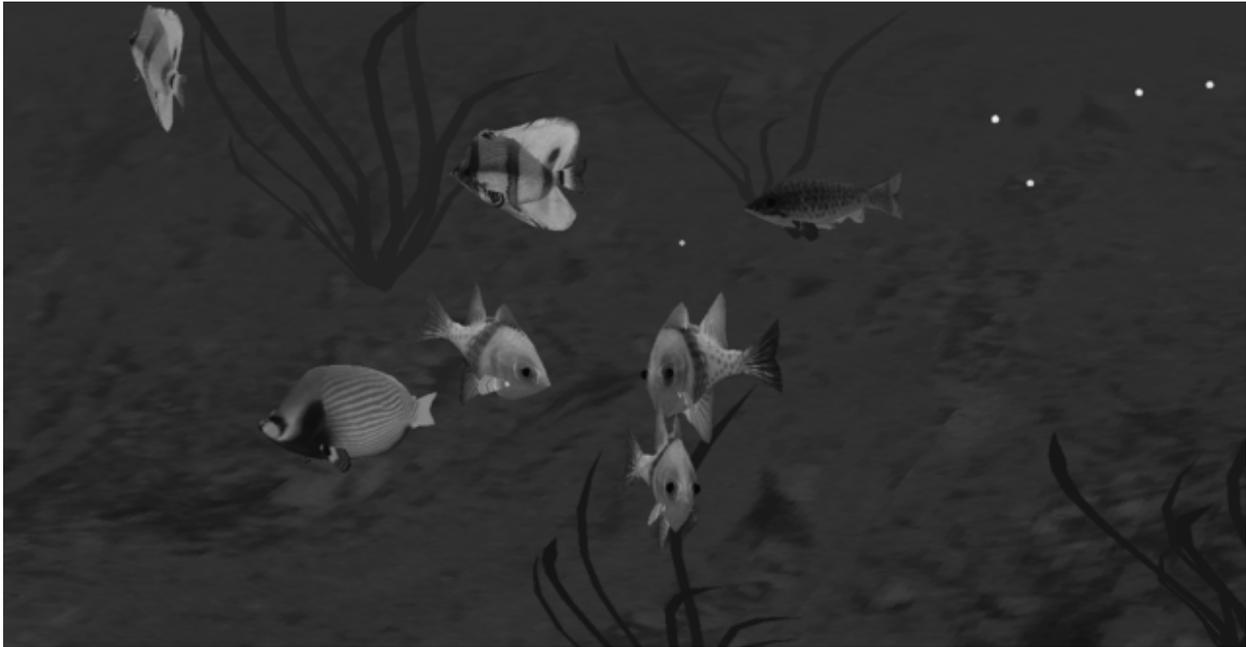


Figure 1.1: Artificial fishes in their physics-based world. See the original color image in Appendix D.



Figure 1.2: Mating behavior. Female (top) is courted by larger male. See the original color image in Appendix D.



Figure 1.3: Predator shark stalking school of prey fish. See the original color image in Appendix D.

animals emulate range from reflexive behaviors to motivational behaviors, and from complex individual behaviors to elaborate group behavior. It is important to appreciate that our goal is *not* to attempt to replicate the complete behavioral repertoire of any one specific fish species but, rather, to develop a generic behavioral model suited to the animation of various species of fishes.

The main contributions of this thesis in more detail are as follows:

- We develop an animation framework that, with minimal intervention from the animator, can achieve the intricacy of motion evident in certain natural ecosystems. This framework encompasses realistic appearance, movement, and behavior of individual animals, as well as the patterns of behavior evident in groups of animals. In addition, unlike many other animation systems, our framework has promise for interactive graphical applications, such as virtual reality. Our paradigm is validated by a physics-based, virtual marine world inhabited by a variety of realistic artificial fishes. In particular, we have developed:
 1. An efficient, physics-based graphical fish model:
 - We introduce the first graphical fish model to yield life-like aquatic motions without keyframing. Our physics-based fish model captures the streamlined shape, the muscular structure, and the general biomechanical properties of natural fishes.
 - We have constructed a set of motor controllers that effectively control the muscles

of the artificial fish to generate realistic fish locomotion.

2. A perception model: The perception model simulates essential visual abilities and limitations. It is equipped with a perceptual attention mechanism which is lacking in previous perception models for animation. This perceptual attention mechanism is essential for the realistic modeling of behavior.
 3. A behavior model:
 - A model of the internal motivations of an animal which comprises the innate characteristics of an animal and its dynamic mental state.
 - A set of behavior routines that implement a range of individual and group piscine behaviors that are common across many species, such as collision avoidance (in the presence of both static and moving obstacles), foraging, wandering, searching for comfortable niches, fleeing, schooling and mating.
 - An “intention generator” that arbitrates among different behaviors and controls the perceptual attention mechanism.
- This thesis provides a new experimental environment for research in related disciplines, such as computer vision and robotics. For example, artificial fishes have been used to design an active computer vision system and evaluate its performance (Terzopoulos and Rabie, 1995) and they have been employed to develop algorithms for learning locomotion and other motor skills (Grzeszczuk and Terzopoulos, 1995). Artificial fishes are virtual robots situated in a continuously dynamic 3D virtual world. They offer a much broader range of perceptual and animation capabilities, lower cost, and higher reliability than can be expected from present-day physical robots used in hardware vision (Terzopoulos, 1995). For at least these reasons, artificial fishes in their dynamic world can serve as a proving ground for theories that profess to account for the sensorimotor competence of animals.

1.4.2 Auxiliary Technical Contributions

- Interactive tools for capturing the realistic appearance of fishes from pictures: We have developed a new tool for efficiently obtaining texture boundaries and coordinates for texture mapping using a deformable mesh.
- A model of a marine environment:
 - We model the physical properties of the hydrodynamic medium in order to simulate its effect on fish motion.

- We have developed a physics-based model of seaweeds that can sway realistically in simulated water current.
- Numerical algorithms for simulating the biomechanical fish models and their physics-based environment: The simulator employs an efficient sparse matrix storage scheme and a fast yet numerically stable semi-implicit Euler method.²
- A graphical user interface:
 - The interface enables the user to create his or her own virtual marine world. For example, the user can decorate the marine environment with static objects and seaweeds, and can specify the number, type, size, initial positions as well as individual habits (i.e. innate behavioral characteristics) of artificial fishes.
 - It allows the user to experiment with the physical properties of the fish, the hydrodynamic medium, and the mental parameters of each fish.
 - It controls the display of a binocular fish-view from the “eyes” of any chosen fish.

1.5 Thesis Overview

The thesis is organized as follows:

In Chapter 2 we review previous work upon which our research draws. At its lowest level of abstraction, our work is an instance of physics-based graphics modeling. Therefore we first survey physics-based modeling where we discuss two basic approaches: the constraint-based approach and the motion synthesis approach. The modeling and control technique we employ in producing realistic fish locomotion pertains to the latter. At a higher level of abstraction, our research is an instance of advanced behavioral animation. We survey prior behavioral animation work and describe related previous perception models developed for the purposes of animation. Then we proceed to discuss the design of action selection mechanisms and review some related previous work in ALife/Animat research.

In subsequent chapters, we describe in detail the animation system that we have developed. In Chapter 3 we begin by presenting a functional overview of the artificial fish model.

²It supports the real-time simulation and wire-frame display (30 frames/second) of up to five swimming fish on a Silicon Graphics R4400 Indigo² Extreme desktop workstation. If real-time performance is not an issue, a huge number of fish may be simulated and rendered photorealistically on such a system.

In Chapter 4 we describe the biomechanical model and how it achieves muscle-based hydrodynamic locomotion. Next, we develop a numerical simulation of the equations of motion. Subsequently, the motor control of the physics-based artificial fish, derived from piscine biomechanical principles, is presented. This includes the construction of the muscle motor controllers as well as the pectoral fin motor controllers.

Chapter 5 describes our approach to constructing geometric display models that capture the form and appearance of a variety of artificial fishes. An interactive, deformable contour tool is developed and applied to map realistic textures over the fish bodies. Finally we explain how the geometric display models are coupled to the biomechanical fish model to yield realistic animation.

In Chapter 6 we present the perception model employed within the artificial fish. In particular, we describe the modeling of the perceptual attention mechanism and the use of motor preferences for generating compromised actions. We present concrete examples of how perception guided behaviors are synthesized. Possible extensions of our perception model are also discussed.

In Chapter 7 we discuss behavior modeling in the artificial fish. We describe the internal motivations and action selection, which is carried out by an intention generator and a set of behavior routines that are explained in detail. Results are presented to illustrate the various behaviors achieved in three varieties of artificial fishes: pacifists, prey, and predators. We then analyze the properties of the action selection mechanism that we have designed and possible extensions are suggested.

Chapter 8 discusses the modeling of the marine environment of the artificial fishes. In particular, we describe the physics-based modeling of seaweeds, food particles and water currents.

In Chapter 9 we present the user interface that we have designed to facilitate the use of our animation system.

In Chapter 10 we describe the animation results that we have achieved to date.

In Chapter 11 we review the contributions of the thesis, and list possible directions of future research.

Chapter 2

Background

In this chapter, we review prior work in the fields of computer graphics and artificial life upon which our research draws. Starting from the lowest level, physics-based graphics modeling, we progressively survey related research on behavioral animation, including perception models for animation and action selection mechanisms. We conclude by putting our work in perspective relative to prior artificial life research on animats.

2.1 Physics-Based Modeling

At its lowest level of abstraction, our work is an instance of physics-based graphics modeling. This approach involves constructing dynamic models of animated objects and computing their motions via physical simulation. Physics-based modeling implies that object motions are governed by the laws of physics, which leads to physically realistic animation. Moreover, this approach frees the animator from having to specify many low-level motion details, since motion is synthesized automatically by the physical simulation. This is evident especially when animating passive motion (i.e. motions of inanimate objects)—the animator need only supply the initial state of the object and a physical simulator automatically computes its motion by integrating the differential equations stemming from Newton's laws.

The success of physics-based modeling was demonstrated in modeling the movements of inanimate objects, such as deformable objects (Terzopoulos et al., 1987; Terzopoulos and Fleischer, 1988; Witkin and Welch, 1990), chains (Barzel and Barr, 1988) and tree leaves (Wejchert and Haumann,

1991). A substantial amount of research has also been concerned with the motion of animate objects, such as humans and animals (Armstrong and Green, 1985; Wilhelms, 1987; Badler, Barsky and Zeltzer, 1991; Hodgins et al., 1995).

An animator requires control over physics-based models in order to produce useful animations. We can categorize physics-based control techniques into two approaches: the *constraint-based approach* and the *motion synthesis approach*.

2.1.1 Constraint-based Approach

The constraint-based approach involves the imposition of kinematic constraints on the motions of an animated object (Platt and Barr, 1988). For example, one may constrain the motion trajectories of certain parts of a model to conform to user specified paths. Two techniques have been used to calculate motions that satisfy constraints: the inverse dynamics technique and the constrained optimization technique.

In inverse dynamics, the motion of an animated body is specified by solving the equations of motion. A set of “constraint forces” (or torques) is computed, which causes the animated body to act in accordance with the given constraints (Isaacs and Cohen, 1987; Barzel and Barr, 1988; Witkin and Welch, 1990). The first two works deal with rigid bodies (for the special case of articulated figures), the third with non-rigid structures. Using inverse dynamics, the resulting motions are physically “correct” (in the sense that the animated body responds to forces in a realistic manner), but they may still look unnatural with respect to any specific form of animal locomotion. For example, when modeling the locomotion of a cat using inverse dynamics, the resulting motion may not resemble that of a cat, but rather that of a robot. This is because, as was discussed in the previous chapter, an animal’s movement not only depends on the Newtonian laws of motion but also is subject to biomechanical principles. In particular, the locomotion of an animal is driven by its muscles, which have limited strength. Therefore, contrary to the assumption of the inverse dynamics technique, a real animal cannot produce arbitrary forces and torques so as to move along any pre-specified path.

The idea behind the constrained optimization technique is to represent motion in the *state-time* space and then define an *objective function* or *performance index* and cast motion control as an optimization problem (Witkin and Kass, 1988; Liu, Gortler and Cohen, 1994). The objective function evaluates the resulting motion. The usual assumption is that motions requiring less energy are preferable. An open-loop controller is synthesized by searching for values of the state space

trajectory, the forces, and the torques that satisfy the constraints and minimize the objective function. This is usually achieved by numerical methods that iteratively refine a user supplied initial guess and that are often computationally expensive. Moreover, it is possible for the motion to be over-constrained, in which case the optimization algorithm is given the responsibility of arbitrating between the user defined constraints and the constraints induced by the laws of physics (Funge, 1995). Unfortunately, this means that the compromise solutions produced may not lead to visually realistic motion, especially if the laws of physics are compromised.

2.1.2 Motion Synthesis Approach

The motion synthesis approach to physics-based control bears greater resemblance to how real animals move. It synthesizes the muscles in natural animals as a set of actuators that are capable of driving the dynamic model of a character to produce locomotion. Unlike inverse dynamics, the motion synthesis approach can take into account the limitations of natural muscles, and unlike constrained optimization, it guarantees that the laws of physics are never violated. It also allows sensors to be incorporated into the animate models which establishes sensorimotor coupling or closed-loop control. This in turn enables an animated character to automatically cope with the richness of its physical environment. Since this approach can emulate natural muscles as actuators, it is able to synthesize various locomotion modes found in real animals by emulating their muscle control patterns. The quality of the results will of course depend upon the fidelity with which the relevant biomechanical structures are modeled. The motion synthesis approach offers less direct animator control than the constraint-based approach. For example, it is almost impossible to produce motions of an animate body that exactly follow some given trajectory, especially for multi-body models, such as human bodies (though humans do experience difficulty when attempting to produce specific trajectories).

Several researchers have successfully applied the motion synthesis approach to animation (Miller, 1988; Terzopoulos and Waters, 1990; Lee, Terzopoulos and Waters, 1993; van de Panne and Fiume, 1993; Ngo and Marks, 1993). The artificial fish model that we develop is inspired by the surprisingly effective model of snake and worm dynamics proposed by Miller (1988) and the face model proposed by Terzopoulos and Waters (1990).¹

An essential physical feature of the bodies of snakes and worms and of human faces is that

¹An early draft of our model was developed based on a fish model that Caroline Houle (a former student at the graphics lab) built for one of her course projects. We would like to acknowledge her contribution to our work.

they are deformable. In both Miller's and Terzopoulos and Waters' works, this feature is efficiently modeled by mass-spring systems. The springs are used to simulate simple muscles that are able to contract by varying their rest lengths. Like these previous models, our fish model is a dynamic mass-spring-damper system with internal contractile muscles that are activated to produce the desired motions. Unlike these previous models, however, we simulate the system using a semi-implicit Euler method which, although computationally more expensive than simple explicit methods, maintains the stability of the simulation over the large dynamic range of forces produced in our simulated aquatic world. Using mass-spring-damper systems, we also model the passive dynamic plants found in the artificial fish habitat.

The major task in motion synthesis is to derive suitable actuator control functions, in particular the time-varying muscle actuator activation functions for different modes of locomotion, such as hopping or flying. When activated according to the corresponding function, each muscle generates forces and torques causing motion of the actuated body parts. The aggregate motion of all body parts forms the particular locomotion pattern. The derivation of actuator control functions becomes increasingly difficult as the number of muscles involved in controlling the locomotion increases. There are two approaches to deriving control functions: the manual construction of controllers and optimization-based controller synthesis, also known as optimal control or learning.

Hand-Crafted Controllers

The manual construction of controllers involves hand crafting the control functions for a set of muscles. This is often possible when the muscle activation patterns in the corresponding real animal are well known. For example, in Miller's (1988) work, the crawling motion of the snake is achieved via a sinusoidal activation (i.e. contraction) function of successive muscle pairs along the body of the snake. Another manually constructed controller for deformable models is that for controlling facial muscles (Terzopoulos and Waters, 1990; Lee, Terzopoulos and Waters, 1995) for realistic human facial animation, where a "facial action coding system" controller coordinates the actions of the major facial muscles to produce meaningful expressions. Most manually constructed controllers have been developed for rigid, articulated figures: Wilhelms (1987) developed "Virya" — one of the earliest human figure animation system that incorporates forward and inverse dynamic simulation; Raibert (1991) showed how useful parameterized controllers of hoppers, kangaroos, bipeds, and quadrupeds can be achieved by decomposing the problem into a set of manually-manageable control problems; Hodgins *et al.* (1995) used similar techniques to animate a variety of human motions

associated with athletics; McKenna *et al.* (1990) produced a dynamic simulation of a walking cockroach controlled by sets of coupled oscillators; Brooks (1991) achieved similar results for a six-legged physical robot; Stewart and Cremer (1992) created a dynamic simulation of a biped walking by defining a finite-state machine that adds and removes constraint equations. A good survey of these sorts of approaches can be found in the book by Badler, Barsky and Zeltzer (1991).

Fish animation poses control challenges characteristic of highly deformable, muscular bodies, not unlike those of snakes (Miller, 1988). We have devised a motor control system that achieves muscle-based, hydrodynamic locomotion by simulating the dynamic interactions between the artificial fish's deformable body and its aquatic environment. To derive the muscle control functions for fish locomotion, we have consulted the literature on marine biomechanics (Webb, 1989; Blake, 1983; Alexander, 1992). The resulting parameterized controllers harness the hydrodynamic forces on fins to achieve forward locomotion over a range of speeds, to execute turns, and to alter body roll, pitch, and yaw so that the fish can move freely within its 3D virtual world.

The main drawback with manually constructed controllers is that they can be extremely difficult and tedious to derive, especially for many-degree-of-freedom body motions. Moreover, the resulting controllers may not be readily transferable to different models or systems; nevertheless, they can serve as a good starting point for the optimization-based algorithms.

Optimization-Based Controller Synthesis

One can also use optimization techniques to derive control functions automatically. An optimization algorithm tries to produce an *optimal controller* through repeated controller and trajectory generation, rewarding better generated motions according to some user specified objective function. This generate-and-test procedure resembles a trial-and-error learning process in humans and animals and is therefore often referred to as "learning". The resulting motion can be influenced indirectly by modifying the objective function. We shall emphasize the difference between the optimization algorithm that we are describing here and the constrained optimization technique mentioned earlier. Here, the laws of physics are not treated as constraints and motion is represented in the *actuator-time* space, rather than in the state-time space.

Since motions are always generated in accordance with the laws of physics, the optimization algorithm is able to exploit the mechanical properties of the physics-based models as well as their environment (Funge, 1995). Interesting modes of locomotion have been automatically discovered

by using simple objective functions that reward low energy expenditure and distance traveled in a fixed time interval (Pandy, Anderson and Hull, 1992; van de Panne and Fiume, 1993; Ngo and Marks, 1993; Sims, 1994; Grzeszczuk and Terzopoulos, 1995). The resulting motions bear a distinct qualitative resemblance to the way that animals with comparable morphologies perform similar locomotion tasks. We shall emphasize, again, that the fidelity of the dynamic model is critical to the realism of the resulting locomotion.

Although we have hand crafted the control functions for the artificial fish’s muscles, our model is rich enough to allow such control functions to be obtained automatically through optimization, as is demonstrated by the work of Grzeszczuk and Terzopoulos (1995).

2.2 Behavioral Animation

At a higher level of abstraction, we are interested in animating the behaviors of animals with an intermediate level of behavioral complexity—somewhere in between the complexity of invertebrates and of primates such as humans. In this regard, our research is an instance of behavioral animation, where the motor actions of characters are controlled by algorithms based on computational models of behavior (Reynolds, 1982; Reynolds, 1987). Consequently, an animator is able to specify motions at a higher level, i.e. the behavior level, as opposed to specifying motion at the locomotion level as is done in physics-based modeling. The animator is therefore concerned with the modeling of individual behaviors. Behavioral animation approaches have been proposed to cope with the complexity of animating anthropomorphic figures (Zeltzer, 1982), animating the synchronized motions of flocks, schools, or herds (Reynolds, 1987) and interactive animation control (Wilhelms, 1990).

The seminal work in behavioral animation is that of Reynolds (1987). Creating vivid animations of flocks of birds or schools of fish using conventional keyframing would require a tremendous amount of effort from an animator. This is because, for example, while the overall motion of birds in a flock is highly coordinated, individual birds have distinct trajectories. In keyframing, the animator would have to script each bird’s motion carefully in each keyframe. By contrast, Reynolds proposed a computational model of aggregate behavior. In his approach, each animated character, called a “boid”, is an independent actor navigating its environment. Each boid has three simple behaviors: *separation*, *alignment* and *cohesion*. A boid decides which behavior to engage in at any given time based on its perception of the local environmental conditions, primarily the location of neighboring boids. The motions of the individual boids are not scripted; rather, the organized flock is an emergent

property of the autonomous interactions between individual boid behaviors.

Although Reynolds' model successfully achieves behavioral realism, it pays little attention to locomotion realism. Because a kinematic model is used to control each boid's locomotion, the resulting motion of individual boids can be visually unrealistic and may not scale well to more elaborate motion. Additionally, the behavioral model is limited by its simplicity. In particular, since the goal is to animate flocking behavior, each boid is capable by design only of behaviors that are useful to flocking. Our artificial fishes are "self-animating" in the sense of Reynolds' work, but unlike his procedural boid actors, they are more elaborate physical models that also have much broader and more complex behavior repertoires.

2.2.1 Perception Modeling

Adaptive behavior is supported by perception of the environment as much as it is by action. It is therefore crucial to model perception in artificial autonomous agents,² including animated animals, humans and physical robots. Reynolds' "boids" maintained flock formations through simple perception of other nearby actors (Reynolds, 1987) while Mataric has demonstrated similar flocking behaviors with physical robots (Mataric, 1994). The roach actor described by McKenna *et al.* (1990) retreated when it sensed danger from a virtual hand. Renault *et al.* (1990) advocate a more extensive form of synthetic vision for behavioral actors, including the automatic computation of internal spatial maps of the world. The virtual humans in Thalmann's work (1995) have simulated simple visual, tactile and auditory sensors that enable them to perform tasks such as following a leader or greeting each other and even playing ball games.

Perception modeling for animation, in general, is very different from that for robotics. In an animation system, the detailed geometry of each scene can always be obtained by interrogating the virtual world model, without extensive sensory information processing. In a robotics system, however, this is not true. In fact, in order for a robot to obtain useful perceptual information, a *visual process* needs to be synthesized. This will include algorithms to infer 3D geometry from images, to identify shapes and to produce appropriate representations of objects, etc. On the other hand, in a typical animation system, since a database of all graphical objects exists and is accessible, the main purpose of modeling perception in animated figures is to enforce behavioral realism. This often only requires that the *perceptual capability* of the animal be modeled, such as the field of view and

²An autonomous agent is an entity in a world that can act or behave on its own without explicit external control. Humans and animals are examples of natural autonomous agents.

occlusion. Reynolds' model of boids (which only consists of simple modeling of limited field of view) is a good example.

Our artificial fishes are currently able to sense their world through simulated visual perception within a deliberately limited field of view. Subject to the natural limitations of occlusion, they can sense lighting patterns, determine distances to objects, and identify objects by inquiring the world model database. They are also equipped with secondary nonvisual modalities, such as the ability to sense the local virtual water temperature. More importantly, unlike previous perception models for animation, the artificial fish's perception induces an attention mechanism. This mechanism allows the fish to train its sensors in a task-specific way as well as to provide other important information for producing convincing behavior. Our model of perception is proven to be effective in generating realistic behaviors of the artificial fish.

2.2.2 Control of Behavior

Behavioral control mechanisms working in conjunction with the perception and the locomotion control mechanisms make our artificial fishes autonomous agents.

To achieve a level of behavioral realism consistent with the locomotional abilities of artificial fishes, it is prudent to consult the ethology literature (Tinbergen, 1950; Lorenz, 1973; McFarland, 1987; Adler, 1975). Tinbergen's landmark studies of the three-spined stickleback highlight the great diversity of piscine behavior, even within a single species. The artificial fishes' behavior repertoires are modeled after natural piscatorial behaviors common across several species. We achieve the nontrivial patterns of behavior (including schooling behaviors as convincing as those demonstrated by Reynolds) in stages. First, we implement primitive reflexive behaviors, such as obstacle avoidance, that tightly couple perception to action (Braitenberg, 1984; Resnick, 1987). Then, through an effective action selection mechanism, the primitive behaviors are combined into motivational behaviors whose activation depends also on the artificial fish's mental state, including hunger, libido, and fear.

As the behavioral repertoire broadens, the issue of action selection becomes crucial. In the next section we survey related prior work on the design of action selection mechanisms for autonomous agents.

2.3 The Modeling of Action Selection

Designing autonomous agents has been one of the major concerns in several fields of research. In software engineering, intelligent computer programs have been developed that can be viewed as autonomous software agents. These programs automatically accomplish various software tasks with little intervention from the user. In robotics, current research has concentrated on developing autonomous mobile robots that can function successfully with little or no human monitoring. In computer graphics, by modeling each animated character as an autonomous agent, complex animations can be produced with minimal intervention from the animator, as is demonstrated in this thesis.

Any autonomous agent will encounter the problem of action selection. The task of action selection (also known as behavior arbitration or behavioral choice) is to determine, from a set of available actions, the most appropriate one based on the agent's internal and external conditions. Designing effective action selection mechanisms is a major endeavor in the design of autonomous agents. To this end, two questions need to be answered first: "what do we mean by an *action*?" and "what do we mean by *the most appropriate* action?".

2.3.1 Defining Action

The term "action" has been used widely in the ethology, psychology and robotics literature, sometimes with quite different meanings. A popular definition of action, in the context of action selection in animals, is given by Tyrrell (1992): an action refers to one of the mutually exclusive entities at the level of *the behavioral final common path* (McFarland and Sibly, 1975). The level of the behavioral final common path is the lowest level of control in an animal or agent's control system, whereby all behaviors are expressed. However meaningful and accurate from an ethologist's point of view, this definition leads to a certain degree of confusion from a designer's point of view. More precisely, it gives the impression that an action is "a movement", hence the action selection mechanism in an anthropomorphic robot, for example, is responsible for determining the detailed movements of every muscle. As a result, one cannot help wondering how the action selection mechanism is any different from the whole control system of an agent. An alternative is therefore to define an action as a *motor skill*, such as "move forward" or "turn to the right".³ Depending on the mechanical model of an agent, each of its actions may require a simple single movement of a single actuator or muscle,

³In fact, although often not explicitly stated, actions are most commonly defined as motor skills in robotics.

or a complex series of coordinated movements of a number of actuators.

Defining an action as a motor skill allows us to deal with the problem of action selection at a higher level by differentiating mechanisms for *motor control* from those for action selection. More specifically, given this definition, motor control mechanisms are responsible for controlling and coordinating actuators in an agent so as to form useful motor skills, i.e. actions, while action selection mechanisms are only responsible for choosing an action without knowing how it is implemented. This concept underlies the design of the artificial fish. We first build a motor control system to implement a set of basic motor skills, including swim forward and backward, turn right and left, glide, yaw, pitch, roll, and brake. The behavior control system is then built to control the selection of these motor skills in order to produce realistic behaviors. Fig. 2.1 illustrates how action selection differs from motor control in a general design scheme. Note that exclusive actions, i.e., actions that use the same actuators/muscles, cannot be selected simultaneously. For example, one can not walk while sitting. However, non-exclusive actions can be selected simultaneously. For example, one can walk while eating.

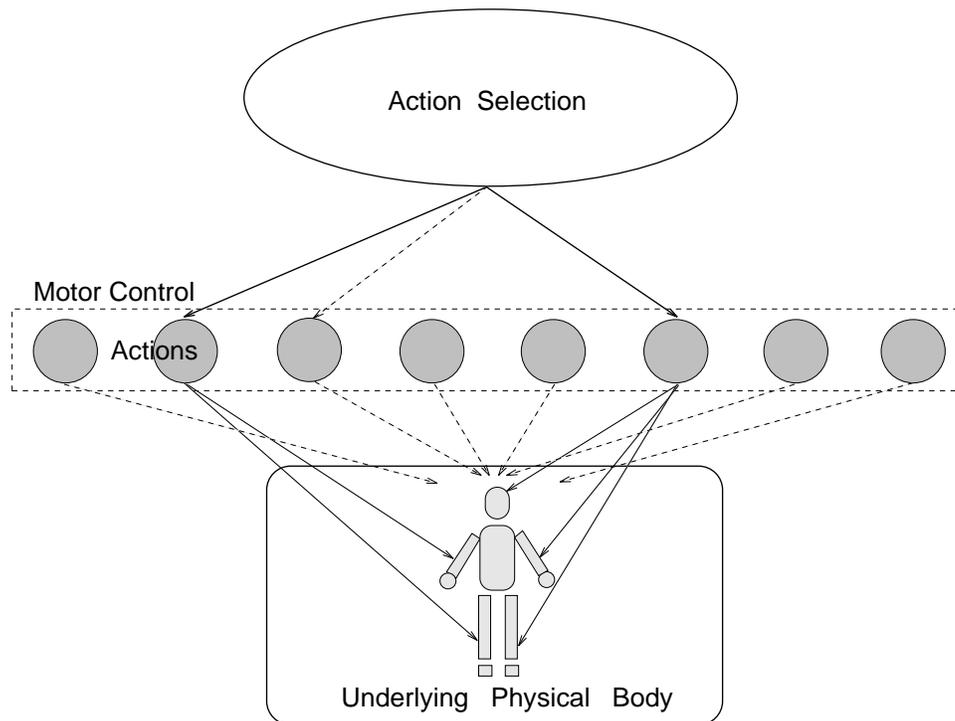


Figure 2.1: Differentiating action selection from motor control in design.

It should be realized that, currently, motor control in complex animated creatures, such as articulated figures, still remains an unsolved problem. This limits the present feasibility of implementing

the above described behavior control scheme in such creatures. However, this should not indicate the lack of generality of the control scheme for we believe that the separation of action selection from motor control is necessary in achieving high level behaviors in most creatures, especially complex ones.

2.3.2 Goals and Means

The most appropriate action can be evaluated with respect to the goal of action selection. Animals in the wild often face hazardous situations. The appropriate choice of actions is crucial to their long term survival. Therefore, according to Dawkins (1989), the ultimate goal of action selection for animals is to choose successive actions (or behaviors) so as to maximize the number of copies of its genes in future generations. That is to say, the ultimate goal is to survive and to reproduce. This goal breaks down to more immediate, day to day behavioral needs. For a robot, the action selection problem entails maintaining the safety of the robot while pursuing the successful completion of the tasks it has been assigned. For a virtual animal (or an autonomous animated creature in general), the goal of action selection is to achieve satisfactory behavioral realism. Since a mathematical representation of the above goals would be extremely complex, we may not be able to build action selection mechanisms via numerical methods, such as optimization. This makes the problem of deriving action selection mechanisms a design problem and hence the main issue is to come up with the corresponding design criteria (Werner, 1994; Maes, 1991a; Tyrrell, 1992).

Although the goal of action selection in a real animal seems different from that in a virtual animal, they are in fact similar. The action selection mechanisms in animals allow them to take appropriate actions in the face of uncertainty. This is what we refer to as rational, adaptive behavior. When we say the behavior of a virtual animal looks realistic, we generally mean that the behavior it exhibits makes sense. Like a real animal, the virtual animal “knows” to avoid hazards, to exploit opportunities and to reasonably allocate resources, etc. Our approach to achieving such realism in animation involves identifying the important principles by which animals select actions, especially priorities between different behaviors, and employing these principles as the design criteria for building action selection mechanisms in virtual animals.

2.3.3 Previous Work

In prior behavioral animation work, such as that of Reynolds (1987), only stimulus-driven action selection process is modeled. For example, the three sub-behaviors of a flocking boid are activated directly by environmental conditions. If the environmental conditions for more than one sub-behavior occur, the one with the highest weight gets chosen (each sub-behavior is associated with a weighting factor that reflects its importance). Comparably, Sun and Green (1993) specify the action selection of a synthetic actor through a set of “relations” each of which is a mapping from an external stimulus to a response of the actor.

Analogous strategies have been taken in robotics. Representative examples include the rule-based mechanism in “Pengi” developed by Agre and Chapman (1987); and reactive systems and systems with emergent functionality, such as those proposed by Kaelbling (1987), Brooks (1986) and Anderson (1990). These mechanisms have the main advantage of coping well with contingencies in the environment since actions are more or less directly coupled with external stimuli. However, as Tyrrell (1993a) points out:

... while we realise that many traditional planning approaches are unsuitable for action selection problems due to their rigidity and disregard of the uncertainty of the environment, we also realize that stimulus-driven mechanisms err in the opposite direction.

In particular, these mechanisms do not model the agent’s internal state thus cannot take into account the agent’s motivations. They are therefore limited in dealing with more sophisticated action selection problems faced by agents with multiple high-level (probably motivational) behaviors, such as those faced by most animals.

Being well aware of the aforementioned problems, researchers in ALife and related fields (e.g. robotics) have come up with various implementation schemes for action selection in animats that takes into account both internal and external stimuli. This body of work provides valuable reference to our design of the behavior control system of the artificial fish.

Maes (1990; 1991a) proposed a distributed, non-hierarchical implementation of action selection, called a “behavior choice network”. The results demonstrate that this model possesses certain properties that are believed to be important to action selection in real animals, such as persistence in behavior, opportunism and satisfactory efficiency. However, while the distributed structure offers good flexibility, it also causes some problems. For example, convergence to a correct choice of

behavior is hard to guarantee. Using a similar architecture to that of Maes' network, Beer and Chiel (1991) proposed a neuroethology-based implementation of an artificial nervous system for simple action selection in a robot insect.

Along the line of Maes' and Beer's work, others have also proposed various different network-type of action selection mechanisms. For instance, Sahota's (1994) mechanism allows behaviors to "bid", and the behavior with the highest bid represents the most appropriate choice.

A common attribute of the above mechanisms (and many others) is the use of a winner-takes-all selection/arbitration process, where the final decision is made exclusively by the winning action or behavior. While this offers highly focussed attention and hence efficiency, it ignores the importance of generating compromised actions. The ability to compromise between different, even conflicting, desires is evident in natural animals. Tyrrell (1993b) emphasized this particular aspect of animal behavior in the implementation of what is known as *free-flow* hierarchies. A free-flow hierarchy implements compromised actions (Rosenblatt and Payton, 1989) within a hierarchical action selection architecture similar to those proposed by early ethologists, such as Tinbergen (1951). The winner is only chosen at the very bottom level of the hierarchy. Simulation results (Tyrrell, 1993a) show that free-flow hierarchies yield favorable choices of action compared to other mechanisms, such as Maes'. A similar scheme is used to design the brains of the pets in Coderre's (1987) *PetWorld*. The main difference between the decision-making (or data-flow) hierarchy in *PetWorld* and a free-flow hierarchy is that, in the former, sensory data flows from the bottom of the hierarchy to the top while in the latter, it flows top down. The major drawback of such an implementation is its high complexity, hence inefficiency, due to the large amount of computations required.

The behavior system of the artificial fish incorporates both stimulus-driven mechanisms and motivation-based mechanisms for action selection. As a result, the fish possesses a level of behavioral capacity to achieve coherence among a number of complex behaviors. In this regard, our work is compatible with the work by Coderre (1987), Maes (1990; 1991a) and Tyrrell (1993b).

Our implementation is similar to that of Tyrrell in that it employs a top-down hierarchical structure and real valued sensory readings, and it can generate compromised actions. Unlike Tyrrell's model, our mechanism employs essentially a winner-takes-all selection process and allows only certain losing behaviors to influence the execution of the winning behavior. This way action selection is carried out much more efficiently than a free-flow hierarchy. Since more than one behaviors influence the selection of the detailed actions taken for accomplishing the chosen behavior, the final choices of actions are preferable to that generated by a conventional winner-takes-all process.

Moreover, the majority of the previous action selection models (including the aforementioned ones) are based on a discrete 2D world which simplifies the problem by greatly restricting legal choice of motor actions. Our model, however, is based on a continuous 3D environment in which the animated animals perform continuous motions.

2.3.4 Task-level Motion Planning

When one attempts to animate advanced animals, such as humans, it becomes necessary to incorporate more abstract action selection mechanisms, such as mechanisms based on reasoning. This approach involves using AI techniques and is termed *task-level motion planning* (a good reference book on this subject is the book by Magnenat-Thalmann and Thalmann (1990)); The most representative animation work along this line is that by Badler and his group who animated a human figure 'Jack' (Badler, Phillips and Webber, 1993). The planner takes as input Jack's initial state and a 3D representation of his world and generates as output a series of actions necessary for accomplishing an assigned task, such as "go and get some ice cream".

While planning ability certainly is one of the most important characteristics in human behavior (and hence is important to model for animations of humans), it is not known as a common feature in most animals lower on the evolutionary ladder than primates. Rather, animal behavior is believed to rest on the more primitive and more fundamental faculty of reactive or adaptive behavior (Tinbergen, 1951; Manning, 1979; McFarland, 1993a). Adaptive behavior enables animals to be autonomous and to survive in uncertain and dynamic environments. Our approach to behavioral animation reflects the adaptiveness of animal behavior. (Note that we are not referring to adaptiveness in the sense of learning nor evolutionary adaptation, but rather, to the ability to select appropriate behaviors according to the perceived situation.) In our approach, we gain high level control through the construction of a model of adaptive behavior where the actions of an artificial animal result from its *active* interaction with the world as guided by its perception.

2.4 Summary

Our approach to developing artificial animals is consistent with the "animat" approach proposed by Wilson (Wilson, 1990). To render our computational model visually convincing, we attempt also to capture, with reasonable fidelity, the appearance and physics of the animal and its world. Artificial

fishes may be viewed as animats of high sophistication. They are autonomous virtual robots situated in a continuous, dynamic 3D virtual world. Their functional design, including motor control, perceptual modeling, and behavioral simulation presents hurdles paralleling those encountered in building physical autonomous agents (see, e.g., the compilation by Maes (1991b)). Previously, the most complex animats were inspired by insects. Brooks (see (Brooks, 1991)) describes a physical insect robot “Genghis”, bristling with sensors, that can locomote over irregular terrain, while Beer develops a virtual counterpart, a cockroach with simple behaviors in a 2D world (Beer, 1990; Beer and Chiel, 1991). Our work tackles animal behavior more complex than those modeled in existing work such as the above. To deal with the broad behavioral repertoire of fishes, we exploit ideas from physics-based graphics modeling, from biomechanics, from behavioral animation, from autonomous agent studies and from ethology.

Chapter 3

Functional Anatomy of an Artificial Fish

As we discussed in the preceding chapter, there are diverse aspects to the realistic modeling of an artificial animal, from superficial appearance to internal functionality. It is helpful to think of the artificial fish model as consisting of three submodels:

1. A *graphical display model* that uses geometry and texture mapping to capture the form and appearance of any specific real fish.
2. A *biomechanical model* that captures the physical and anatomical structure of the fish's body, including its muscle actuators, and simulates its deformation and physical dynamics.
3. A *brain model* that is responsible for motor control, perception control and behavior control of the fish.

Each of the three submodels will be developed in detail in the coming chapters. Fig. 3.1 presents a functional overview of the artificial fish. As the figure illustrates, the body of the fish harbors its brain. The brain itself consists of three control centers: *the motor center*, *the perception center*, and *the behavior center*. These centers are part of the motor, perception, and behavior control systems of the artificial fish. The function of each of these systems will be previewed next.

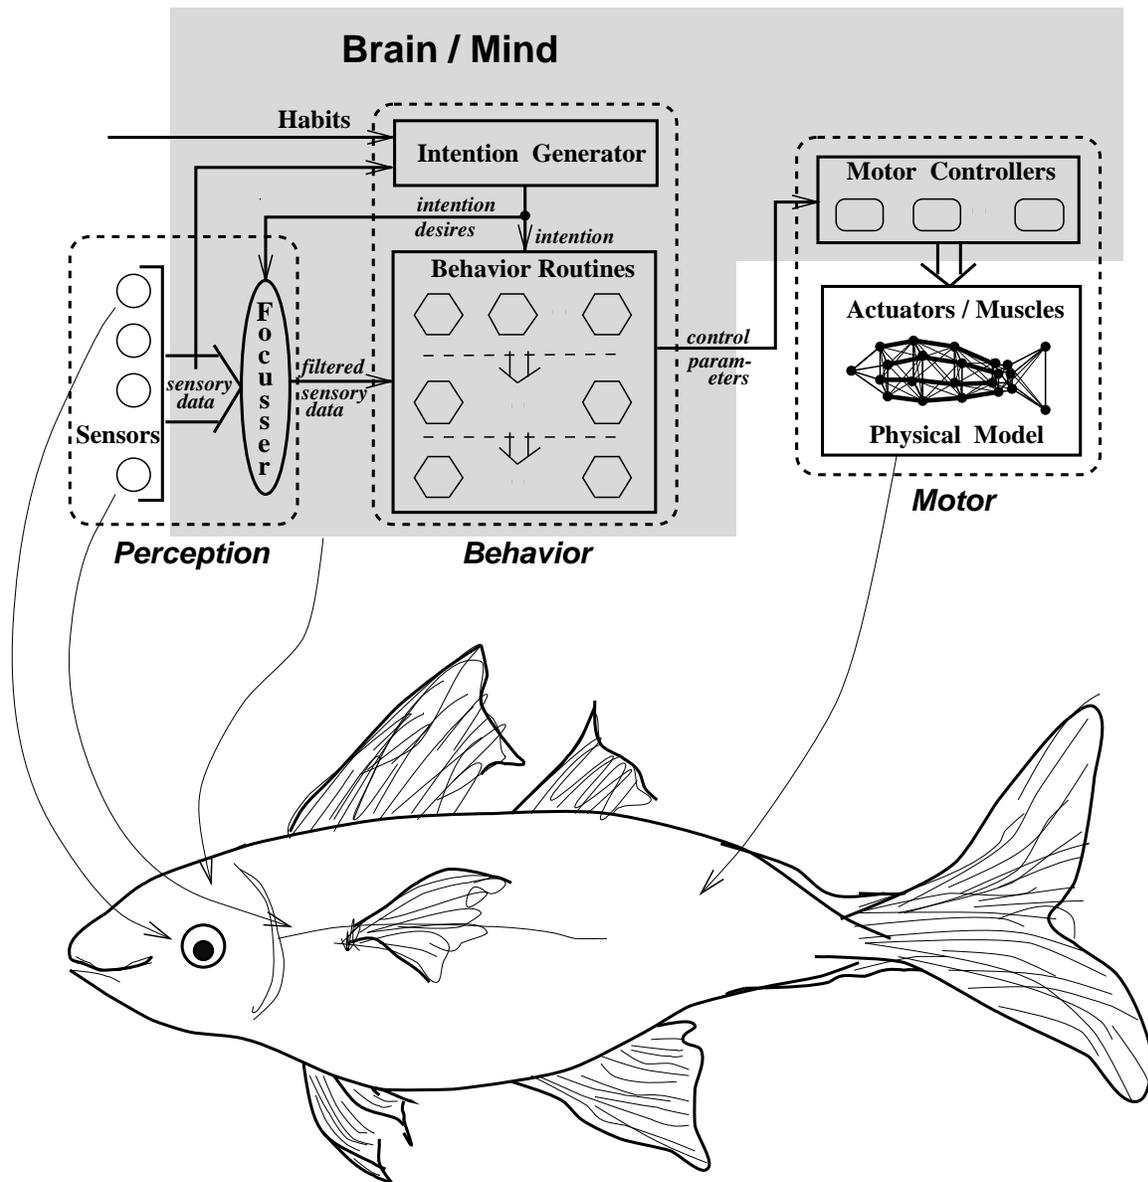


Figure 3.1: System overview of the artificial fish.