

Artificial Fishes: Autonomous Locomotion, Perception, Behavior, and Learning in a Simulated Physical World

Demetri Terzopoulos, Xiaoyuan Tu, and Radek Grzeszczuk

Department of Computer Science, University of Toronto
10 King's College Road, Toronto, Ontario, M5S 1A4, Canada
e-mail: {dt|tu|radek}@cs.toronto.edu

Published in *Artificial Life*, **1**(4):327–351, 1994.

Abstract

This paper develops artificial life patterned after animals as evolved as those in the superclass Pisces. It demonstrates a virtual marine world inhabited by realistic artificial fishes. Our algorithms emulate not only the appearance, movement, and behavior of individual animals, but also the complex group behaviors evident in many aquatic ecosystems. We model each animal holistically. An artificial fish is an autonomous agent situated in a simulated physical world. The agent has (i) a three-dimensional body with internal muscle actuators and functional fins, which deforms and locomotes in accordance with biomechanic and hydrodynamic principles, (ii) sensors, including eyes that can image the environment, and (iii) a brain with motor, perception, behavior, and learning centers. Artificial fishes exhibit a repertoire of piscine behaviors that rely on their perceptual awareness of their dynamic habitat. Individual and emergent collective behaviors include caudal and pectoral locomotion, collision avoidance, foraging, preying, schooling, and mating. Furthermore, artificial fishes can learn how to locomote through practice and sensory reinforcement. Their motor learning algorithms discover muscle controllers that produce efficient hydrodynamic locomotion. The learning algorithms also enable artificial fishes to train themselves to accomplish higher level, perceptually guided motor tasks, such as maneuvering to reach a visible target.

Keywords: artificial life, autonomous agents, animats, artificial fishes, learning, behavior, perception, locomotion, physics-based modeling, computer graphics.

1 Introduction

Imagine a virtual marine world inhabited by a variety of realistic fishes.¹ In the presence of underwater currents, the fishes employ their muscles and fins to swim gracefully around immobile 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.

The computer emulation of the above scenario presents a formidable challenge in the field of artificial life. In this paper we propose a computational framework for creating *fully functional artificial animals*—in this instance, *artificial fishes*. Artificial fishes are autonomous agents with functional bodies controlled by brains. Their appearance, motivations, and complicated group interactions aspire to be as faithful as possible to nature's own. To simulate artificial worlds with the level of complexity of the one depicted above, we have taken a bottom-up, compositional approach. In our approach we model not just 3D form and appearance, but also the basic physics of the animal and its environment. Upon the simulated physics substrate, we can effectively model the animal's means of locomotion. This in turn positions us to model the animal's perceptual awareness of its world, its behavior, and its ability to learn. Our holistic approach to modeling the animal and its world is crucial to achieving realism.

The long-term goal of our research is a computational theory that can potentially account for the interplay of physics, locomotion, perception, behavior, and learning in higher animals. A good touchstone of such a theory is its ability to produce visually convincing results in the form of realistic computer graphics animation with little or no animator intervention. We have been able to achieve such results in two instances to date. Our animation "Go Fish!" [23] 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." Plate 1a shows a variety of animated artificial fishes. The reddish fish are engaged in a mating ritual, the greenish 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 Plate 1b, 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. Plate 1c 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"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 [25].



Plate 1a: Artificial fishes in their physics-based world.



Plate 1b: Mating behaviors. Female (top) is courted by large male.



Plate 1c: Predator shark stalking school of prey fish.

1.1 Background

Our approach to developing artificial animals is consistent with the “animat” approach proposed by Wilson [26]. To render our computational model visually convincing, we attempt as well to capture with reasonable fidelity the appearance and physics of the animal and its world. Artificial fishes may be viewed as animats of unprecedented sophistication. They are autonomous virtual robots situated in a continuously 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 [13]). Previously, the most complex animats were inspired by insects. Brooks (see [8]) 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 [5] (see also [13]).

Our work tackles animals much more highly evolved and complex than insects. To deal with the broad behavioral repertoire of fishes, we exploit ideas from classical ethology [22, 12, 15, 1]. Tinbergen’s landmark studies of the three-spined stickleback highlight the great diversity of piscine behavior, even within a single species. We achieve the nontrivial patterns of behavior outlined in the introductory paragraph of this paper in stages. First, we implement primitive reflexive behaviors, such as obstacle avoidance, that directly couple perception to action [7]. Then we combine the primitive behaviors into motivational behaviors whose activation depends also on the artificial fish’s mental state, including hunger, libido, and fear.

Useful behavior is supported by perception of the environment as much as it is by action. Reynolds’ “boids” maintained flocking formations through perception of other nearby boids [20]. Recently Mataric has demonstrated similar flocking behaviors with physical robots [14]. Artificial fishes sense their world through simulated visual perception within a deliberately limited field of view. They can sense lighting patterns, determine distances to objects, and identify objects, subject to the natural limitations of occlusion. Furthermore, they are equipped with secondary nonvisual modalities, such as the ability to sense local water temperature.

At its lowest level, our work makes use of computational physics. We model the biomechanics of a broad class of fishes and their muscle-based locomotion abilities that exploit the physics of their liquid medium [6, 2]. The mechanical model that we develop was inspired by the simple but surprisingly effective computer graphics model of snake and worm dynamics proposed by Miller [17].

We have provided artificial fish with algorithms that enable them to learn automatically from first principles how to

achieve hydrodynamic locomotion by controlling their internal muscle actuators. The locomotion learning algorithm that we describe is more continuous and closer connected to actuation than most of the animat “behavior learning” algorithms surveyed in [16]. Our multilevel reinforcement learning procedure first performs a global search for actuator activation functions that produce efficient locomotion. The process then abstracts these activation functions into a highly compact representation. The representation emphasizes the natural periodicities of the derived muscle actions and makes explicit the coordination among multiple muscles that leads to effective locomotion. Finally, the artificial fish can put into practice the compact, efficient controllers that it has learned and train itself to accomplish higher level sensorimotor tasks.

1.2 Functional Overview of the Artificial Fish

Fig. 1 presents an overview of an artificial fish situated in its simulated physical world. The body of the fish harbors a brain or mind with motor, perception, behavior, and learning centers.




Figure 1: Control and information flow in artificial fish.

The motor system, comprising the actuators and a set of motor controllers (MCs), drives the dynamic model of the fish. We have crafted a mechanical body model that represents a good compromise between anatomical consistency, hence realism, and computational efficiency. Our model is rich enough so that we can build MCs by gleaning information from the fish biomechanics literature [6, 2]. The MCs are parameterized procedures. Each is dedicated to carrying out a specific motor function, such as “swim forward” or “turn left.” They translate natural control parameters such as the forward speed or angle of the turn into detailed muscle actions.

The perception system relies on a set of on-board virtual sensors to provide sensory information about the dynamic environment, including eyes that can produce time-varying retinal images of the environment. The brain’s perception center includes a perceptual attention mechanism which allows the artificial fish to train its sensors at the

world in a task-specific way, hence filtering out sensory information superfluous to its current behavioral needs. For example, the artificial fish attends to sensory information about nearby food sources when foraging.

The behavior center of the artificial fish’s mind mediates between its perception system and its motor system. An intention generator, the fish’s cognitive faculty, harnesses the dynamics of the perception-action cycle. The innate character of the fish is established by a set of habits that determine whether or not it is male or female, predator or prey, etc. The intention generator combines the habits and mental state with the incoming stream of sensory information to generate dynamic goals for the fish, such as to hunt and feed on prey. It ensures that goals have some persistence by exploiting a single-item memory. The intention generator also controls the perceptual attention mechanism. At every simulation time step, the intention generator activates behavior routines that attend to sensory information and compute the appropriate motor control parameters to carry the fish one step closer to fulfilling its current intention. Primitive behavior routines, such as obstacle avoidance, and more sophisticated motivational behavior routines, such as mating, implement the behavioral repertoire of the artificial fish.

The learning center of its mind enables the artificial fish to learn how to locomote through practice and sensory reinforcement. Through optimization, the motor learning algorithms discover muscle controllers that produce efficient locomotion. Their brain’s learning center also enable artificial fishes to train themselves to accomplish higher level sensorimotor tasks, such as maneuvering to reach a visible target.

2 Realistic Modeling of Form and Appearance

To achieve realism, our artificial fish model must first represent the form and appearance of real fishes with sufficient fidelity. To this end, we have perused photographs of real fishes, such as those shown in Fig. 2(a), and have built 3D geometric models of several different species using nonuniform rational B-spline (NURBS) surfaces (Fig. 2(b)).

The next step is to map realistic textures onto the geometric fish model (Fig. 2(e)). We extract natural textures from digital images of the fish photos, employing a “snake-grid” tool to determine appropriate texture map coordinates in the image. Snakes are interactive deformable contours that are subject to a force field derived from an image [11]. The force field attracts them towards interesting image features such as intensity edges. A collection of coupled snakes forms a deformable grid (Fig. 2(c–d)). The snake-grid floats freely over an image and it can be pulled into position using the mouse. When its border approaches the intensity edges that demarcate the fish from its background in the image, the border snakes lock on and adhere to these edges. The remaining snakes in the grid relax elastically to cover the imaged fish body with a nonuniform but smooth coordinate system (Fig. 2(d)). The snake crossing points serve as texture map image coordinates for the NURBS surfaces.

3 Physics-Based Fish Model and Locomotion

Studies into the dynamics of fish locomotion show that most fishes use their caudal fin as the primary motivator [24]. Caudal swimming normally uses posterior muscles on either side of the body, while turning normally uses anterior muscles. To synthesize realistic fish locomotion we have designed a dynamic fish model consisting of 23 nodal point masses and 91 springs. The spring arrangement maintains the structural stability of the body while allowing it to flex. Twelve of the springs running the length of the body also serve as simple muscles (Fig. 3).

3.1 Mechanics

The mechanics of the spring-mass model are specified as follows: Let node i have mass m_i , position $\mathbf{x}_i(t) = [x_i(t), y_i(t), z_i(t)]$, velocity $\mathbf{v}_i(t) = d\mathbf{x}_i/dt$, and acceleration $\mathbf{a}_i(t) = d^2\mathbf{x}_i/dt^2$. Let elastic spring S_{ij} connect node i to node j and denote its spring constant as c_{ij} and natural, rest length as l_{ij} . Its deformation is $e_{ij}(t) = \|\mathbf{r}_{ij}\| - l_{ij}$, where $\mathbf{r}_{ij}(t) = \mathbf{x}_j - \mathbf{x}_i$. The force S_{ij} exerts on node i is $\mathbf{f}_{ij}^s(\mathbf{x}_i, \mathbf{x}_j) = c_{ij}e_{ij}\mathbf{r}_{ij}/\|\mathbf{r}_{ij}\|$ (note that $\mathbf{f}_{ij}^s = -\mathbf{f}_{ji}^s$). The Lagrange equations of motion of the dynamic fish are:

$$m_i \frac{d^2\mathbf{x}_i}{dt^2} + \rho_i \frac{d\mathbf{x}_i}{dt} - \mathbf{w}_i = \mathbf{f}_i^w; \quad i = 0, \dots, 22, \quad (1)$$

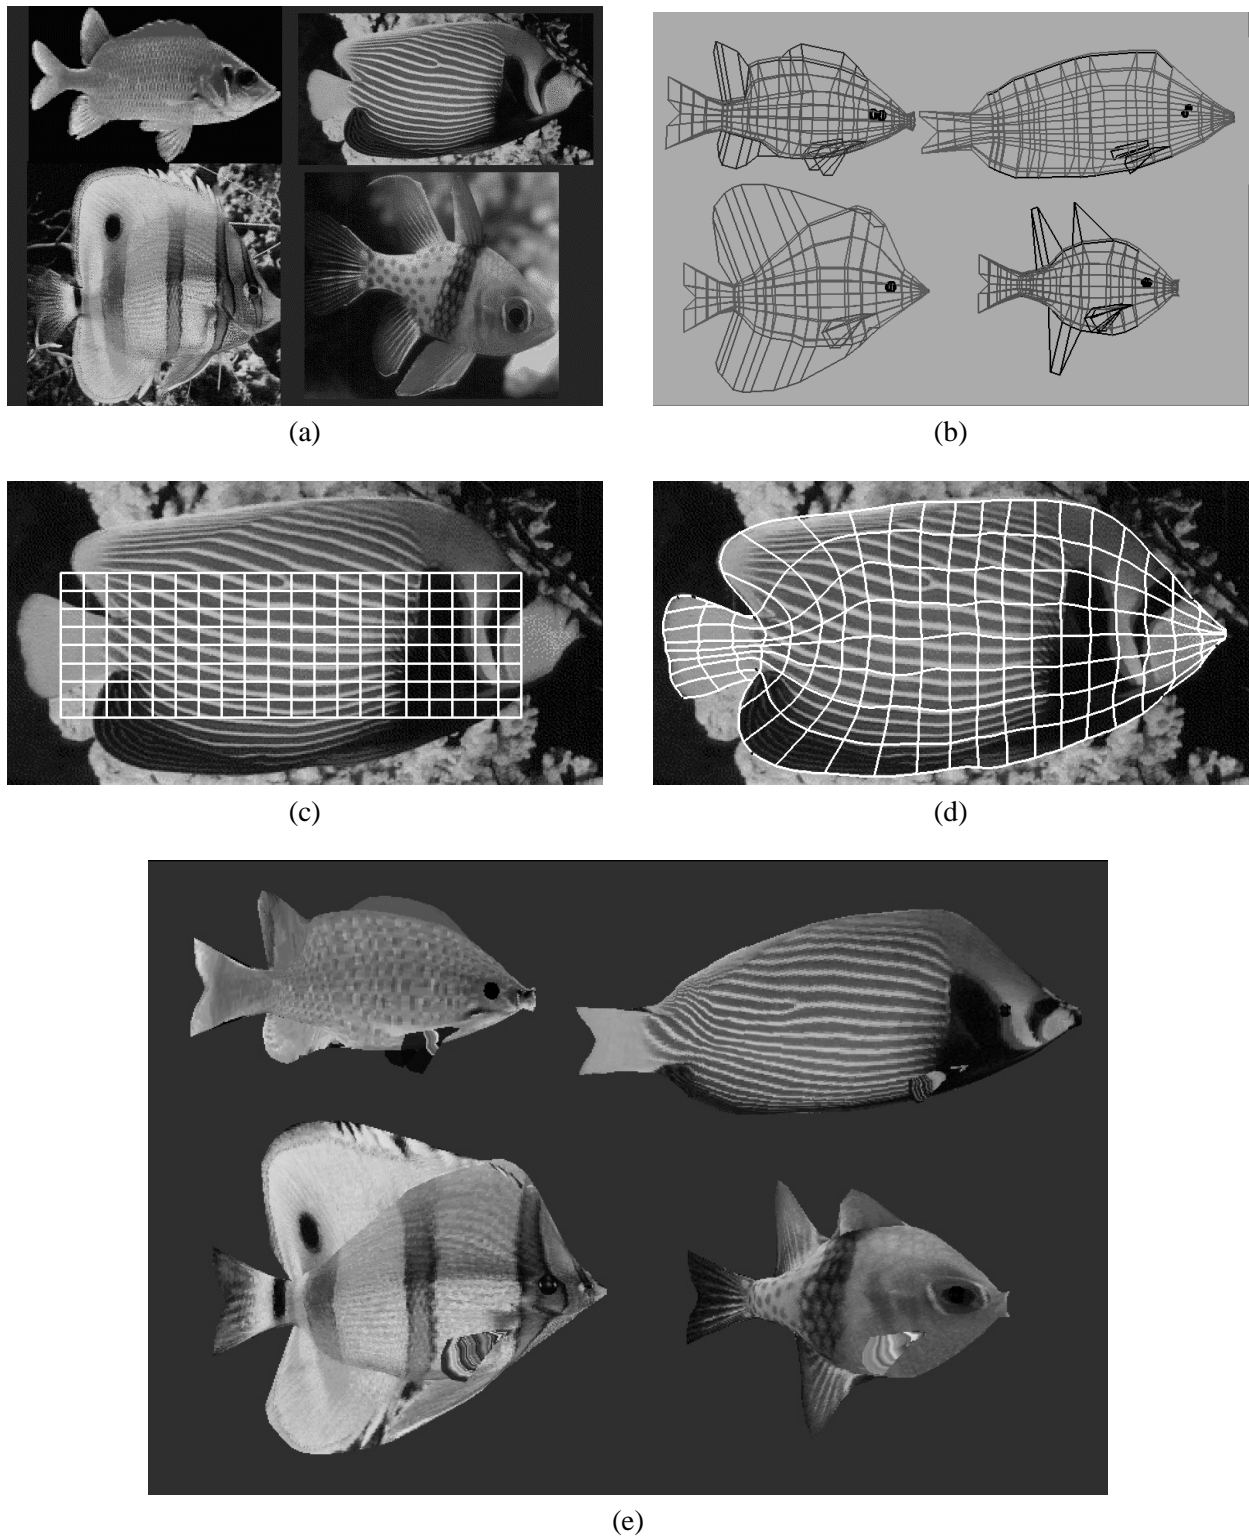


Figure 2: (a) Digitized images of fish photos. (b) 3D NURBS surface fish bodies. Initial (c) and final (d) snake-grid covering an imaged fish body. (e) Texture mapped 3D fish models.

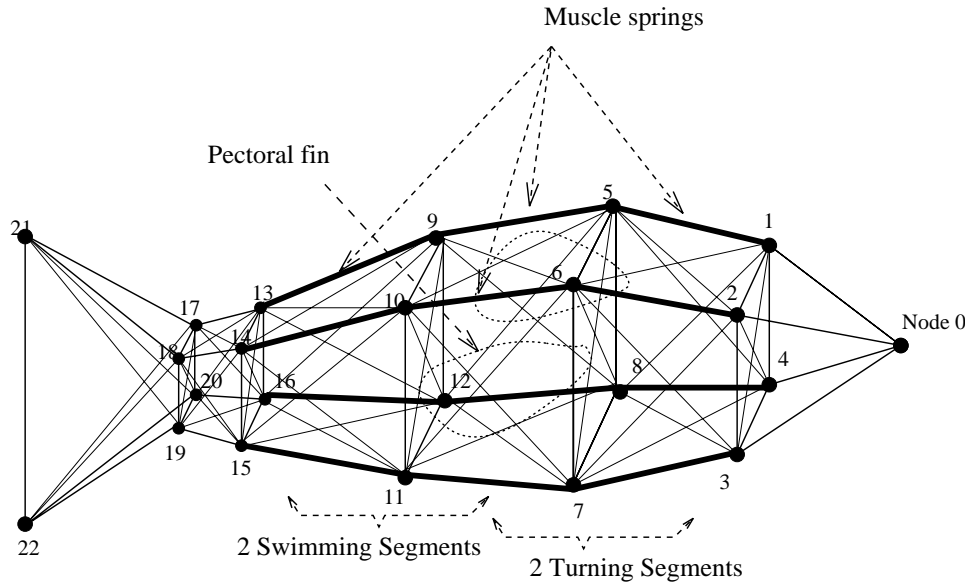


Figure 3: Dynamic fish model. Nodes are lumped masses. Lines are springs (shown at their natural lengths). Bold lines are muscle springs.

where ρ_i is the damping factor, $\mathbf{w}_i(t) = \sum_{j \in N_i} \mathbf{f}_{ij}^s$ is the net internal force on node i due to springs connecting it to nodes $j \in N_i$, where N_i is the index set of neighboring nodes. Finally, $\mathbf{f}_i^w(t)$ is the external (hydrodynamic) force on node i .

To integrate the differential equations of motion, we employ a numerically stable, implicit Euler time stepping method [19]. Since the elastic forces depend nonlinearly on the position variables \mathbf{x}_i , the method assembles the sparse stiffness matrix for the spring-mass system in efficient “skyline” storage format, then factorizes and solves the resulting system of algebraic equations at every time step to obtain the position increments.²

We couple the control points of the aforementioned texture mapped NURBS body model to the time-varying positions of the mass points (the nodes in Fig. 3), such that the fish body deforms in accordance with the simulated dynamics of the actuated spring-mass system.

3.2 Swimming Using Muscles and Hydrodynamics

The artificial fish moves as a real fish does, by contracting its muscles. If S_{ij} is a muscle spring, it is contracted by decreasing the rest length l_{ij} . For convenience, we assign a minimum contraction length l_{ij}^{\min} to the muscle spring and express the contraction factor as a number in the range $[0, 1]$. The brain of the fish controls the muscles continuously through time by specifying the vector of muscle actuation functions $\mathbf{u}(t) = [u_1(t), \dots, u_{12}(t)]$, whose components specify a time-varying contraction factor for each of the 12 muscles. The characteristic undulation of the fish’s tail can be achieved by periodically contracting the swimming segment springs on one side of the body while relaxing their counterparts on the other side.

When the fish’s tail beats, it sets in motion a volume of water. The inertia of the displaced water produces a reaction force normal to the fish’s body proportional to the volume of water displaced per unit time, which propels the fish forward (Fig. 4). Under certain assumptions, the instantaneous force on the surface S of a body due to a viscous fluid is approximately proportional to $-\int_S (\mathbf{n} \cdot \mathbf{v}) \mathbf{n} dS$, where \mathbf{n} is the unit outward normal function over the surface and \mathbf{v} is the relative velocity function between the surface and the fluid. For efficiency, we triangulate the surface of the dynamic fish model between the nodes and approximate the force on each planar triangle as

²In our simulation: $m_i = 1.1$ for $i = 0$ and $13 \leq i \leq 19$; $m_i = 6.6$ for $1 \leq i \leq 4$ and $9 \leq i \leq 12$; $m_i = 11.0$ for $5 \leq i \leq 8$, and $m_i = 0.165$ for $i = 21, 22$. The cross springs (e.g., c_{27}) which resist shearing have spring constants $c_{ij} = 38.0$. The muscle springs (e.g., c_{26}) have spring constants $c_{ij} = 28.0$, and $c_{ij} = 30$ for the remaining springs. The damping factor $\rho_i = 0.05$ and the time step used in the Euler time-integration procedure is 0.055.

$\mathbf{f} = \min[0, -A(\mathbf{n} \cdot \mathbf{v})\mathbf{n}]$, where A is the area of the triangle and \mathbf{v} is its velocity relative to the water. The \mathbf{f}_i^w variables at each of the three nodes defining the triangle are incremented by $\mathbf{f}/3$.




Figure 4: Hydrodynamic locomotion. With tail swinging towards positive X axis, reaction force \mathbf{F}_i^w at point n_i acts along the inward normal. Component \mathbf{F}_{ix}^w resists the lateral movement, while \mathbf{F}_{iy}^w is forward thrust. Aggregate thrust propels fish towards positive Y axis.

3.3 Motor Controllers

Currently the artificial fish has three MCs. The swim-MC produces straight swimming, while the left-turn-MC and right-turn-MC execute turns. The MCs prescribe muscle contractions to the mechanical model. The swim-MC controls the swimming segment muscles (see Fig. 3), while the turn MCs control the turning segment muscles.

According to [24], the swimming speed of most fishes is roughly proportional to the amplitude and frequency of the periodic lateral oscillation of the tail, below certain threshold values. Our experiments with the mechanical model agree well with these observations. Both the swimming speed and the turn angle of the fish model are approximately proportional to the contraction amplitudes and frequencies/rates of the muscle springs.

The swim-MC (swim-MC(speed) $\mapsto \{r_1, s_1, r_2, s_2\}$) converts a swim speed parameter into contraction amplitude and frequency control parameters for the anterior (r_1, s_1) and posterior (r_2, s_2) swim segments. One pair of parameters suffices to control each of the two swim segments due to symmetry—the four muscle springs have identical rest lengths and minimum contraction lengths, identical spring constants, and the contractions of the muscle spring pairs on opposite sides are exactly out of phase. Moreover, the swim-MC produces periodic muscle contractions in the posterior swim segment which lag 180 degrees behind those of the anterior swim segment; hence the mechanical model displays a sinusoidal body shape as the fish swims (see [24]).

By experimenting, we have found a set of four maximal parameters, $\hat{r}_1, \hat{s}_1, \hat{r}_2$ and \hat{s}_2 , which produce the fastest swimming speed. The swim-MC generates slower swim speeds by specifying parameters that have values between 0 and the maximal parameters. For example, $\{0.8\hat{r}_1, \hat{s}_1, 0.7\hat{r}_2, \hat{s}_2\}$ results in a slower-swimming fish.

As mentioned earlier, most fishes use their anterior muscles for turning, and the turn angle is approximately proportional to the degree and speed of the anterior bend, up to the limit of the fish's physical strength [24]. The artificial fish turns by contracting and expanding the springs of the turning segments (Fig. 3) in similar fashion. For example, a left turn is achieved by quickly contracting the left side springs of the segments and relaxing those on the right side. This deflects the fish's momentum and brings it into the desired orientation. Then the contracted springs

are restored to their rest lengths at a slower rate, so that the fish regains its original shape with minimal further change in orientation.

Similarly, the left and right turn MCs ($\text{turn-MC}(\text{angle}) \mapsto \{r_0, s_0, r_1, s_1\}$) convert a turn angle to control parameters for the anterior and posterior turning segments to execute the turn (note that the posterior turning segment also serves as the anterior swim segment). Through experimentation, we established 4 sets of parameter values $P_i = \{r_0^i, s_0^i, r_1^i, s_1^i\}$ which enable the fish to execute natural looking turns of approximately 30, 45, 60, and 90 degrees. By interpolating the key parameters, we define a steering map that allows the fish to generate turns of approximately any angle up to 90 degrees. Turns greater than 90 degrees are composed as sequential turns of lesser angles.

3.4 Pectoral Fins

On most fishes the pectoral fins control pitching (the up-and-down motion of the body) and yawing (the side-to-side motion). The pectorals can be held close to the body to increase speed by reducing drag or they can be extended to increase drag and serve as a brake [25]. Many reef fishes use a pectoral swimming style, keeping their bodies still and using their pectoral fins like oars to achieve fine motion control, including reverse motions, when foraging.

The artificial fish is neutrally buoyant in the virtual water and has a pair of pectoral fins which enable it to navigate freely in its 3D world. The pectoral fins function in a similar, albeit simplified, manner to those on real fishes. For our purposes the detailed movement of the pectoral fins is of lesser interest than the movement of the fish body. To simplify the fish model and its numerical solution, we do not simulate the elasticity and dynamics of the pectoral fins. However, we do approximate the dynamic forces that the pectoral fins exert on the body of the fish to control locomotion.




Figure 5: The pectoral fins

The pectoral fins (Fig. 5) work by applying reaction forces to nodes in the midsection, i.e. nodes $1 \leq i \leq 12$ (see Fig. 3). The fins are analogous to the airfoils of an airplane. Pitch, yaw, and roll control stems from changing their orientations relative to the body; i.e., the angle $\pi/4 \leq \gamma \leq \pi$. Assuming that a fin has an area A , surface normal \mathbf{n} and the fish has a velocity \mathbf{v} relative to the water (Fig. 5), the fin force is $F_f = -A(\mathbf{n} \cdot \mathbf{v})\mathbf{n} = -A(\|\mathbf{v}\| \cos \gamma)\mathbf{n}$ which is distributed equally to the 12 midsection nodes. When the leading edge of a fin is elevated, a lift force is imparted on the body and the fish ascends, and when it is depressed a downward force is exerted and the fish descends. When the fin angles differ the fish yaws and rolls. The artificial fish can produce a braking effect by angling its fins to decrease its forward speed (i.e. $\gamma = \pi$). This motion control is useful in maintaining schooling patterns, for instance.

4 Learning Muscle-Based Locomotion

We have discussed how locomotion controllers may be carefully hand crafted using knowledge gleaned from the piscine biomechanics literature and long hours of experimentation with our mechanical fish model. In this section, we consider the following general question: Given a physics-based model of an animal with internal muscle actuators capable of producing locomotion, such as the fish model of Fig. 3, is it possible for the model to learn from first principles how to control its actuators in order to locomote in a natural fashion? Furthermore, can it employ the controllers that it has learned in order to accomplish higher level tasks guided by sensory perception? We demonstrate affirmative answers to both questions by developing a learning center in the mind of the artificial fish that applies a form of reinforcement learning.

4.1 Learning Strategy

We formulate a two-phase, bottom-up strategy for learning muscle controllers. The artificial fish has a fully functional body, but at first it is “brain-dead” and does not know how to use its muscles to locomote. In phase one, it repeatedly practices a variety of muscle activation functions and remembers activation patterns that improve its locomotion, thus learning how to locomote with increasingly better efficiency. Repeated improvements eventually lead to natural looking locomotion patterns that are optimally efficient.

When an adequate degree of optimization has been achieved in the low-level learning phase, the learning algorithm enters phase two, where it abstracts the activation functions into a highly compact representation. The representation drastically reduces the dimensionality of the learning problem by using basis functions that make explicit the natural periodicities of the derived muscle actions and the pattern of coordination among multiple muscles that yields effective locomotion. The artificial fish trains itself with routine locomotion drills and associates the low-level muscle activation functions that it has learned with specific higher-level tasks that it needs to perform. Finally, it can put to use the compact, efficient controllers that it has learned, to accomplish higher level sensorimotor tasks—for example, it can locomote and execute turning maneuvers to reach a visible target. The learned controllers can subsequently be employed during advanced behaviors such as hunting.

4.2 Low-Level Learning

At the foundation of our approach lies the notion that natural motion patterns are energetically efficient. This allows us to reduce the problem of learning realistic locomotion into a problem of optimizing an objective function, for which various solution techniques are available.

Fig. 6 illustrates the learning algorithm. The objective function takes the form

$$E(\mathbf{u}(t)) = \int_{t_0}^{t_1} (\mu_1 E_u(\mathbf{u}(t)) + \mu_2 E_v(\mathbf{v}(t))) dt, \quad (2)$$

a weighted sum, with weighting variables μ_1, μ_2 , of a term E_u that evaluates the vector of muscle actuator control functions $\mathbf{u}(t)$, which dictate the lengths l_{ij} of muscle springs in the dynamic model (see Sec. 3.2), and a term E_v that evaluates the resulting trajectory $\mathbf{v}(t)$ of the artificial fish. Note that to compute $\mathbf{v}(t)$ and hence E , we must perform a forward simulation of the dynamic model over a time interval $t_0 \leq t \leq t_1$ with the actuation function inputs $\mathbf{u}(t)$.

The term E_u guides the optimization by discouraging large, rapid fluctuations of \mathbf{u} . The rationale is that chaotic muscle actuations usually produce energy-inefficient body motions. We encourage smoothness through the function

$$E_u = -\frac{1}{2} \left(\nu_1 \left| \frac{d\mathbf{u}}{dt} \right|^2 + \nu_2 \left| \frac{d^2\mathbf{u}}{dt^2} \right|^2 \right), \quad (3)$$

with weighting factors ν_1 and ν_2 . The two terms are potential energy densities of linear and cubic splines in time, respectively. The former penalizes muscle effort more than the latter.

The criterion E_v for a good trajectory that we used most often in our learning experiments was the final distance of the fish from a target location. Depending on the task, other possible criteria are: the closeness to a specified speed, the closeness to a specified trajectory, etc.

Learning low level control involves the application of simulated annealing to optimize (2) [19]. Simulated annealing is applied after discretizing the actuator control functions $\mathbf{u}(t)$ by sampling them in the time interval under consideration to obtain the set of discrete samples $\mathbf{u}_i = \mathbf{u}(t_i)$ for $1 \leq i \leq N$ (typically, we set N to 15 time samples, and recover the continuous $\mathbf{u}(t)$ through linear or cubic spline interpolation of the \mathbf{u}_i). The annealing algorithm repeatedly perturbs the \mathbf{u}_i to modify the actuator activation functions that control the muscles in the fish. It retains those perturbations that produce increasingly better locomotion as measured by the objective function E , and sometimes accepts those that don't to escape local minima. Note that after performing a forward simulation using $\mathbf{u}(t)$, the artificial fish can evaluate E with its on-board sensors. Hence, learning proceeds autonomously.

Fig. 7 shows six artificial leopard sharks in a race. The furthest shark has completed only 90 annealing steps,

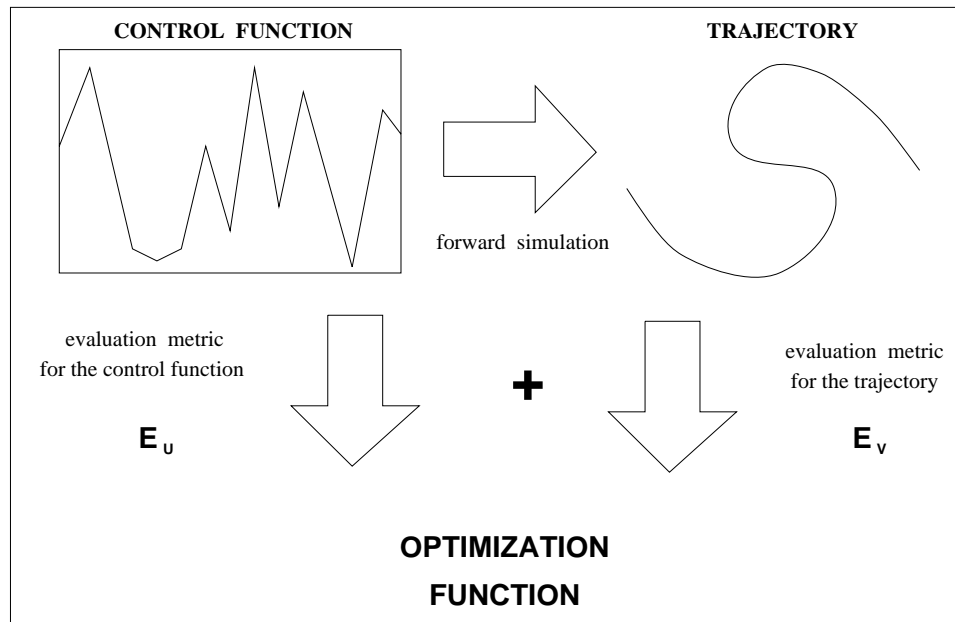


Figure 6: Low-level algorithm for learning locomotion. Muscle control inputs actuate body of artificial animal to produce a locomotion trajectory through forward physical simulation. A fitness function measures the quality of input control and output trajectory.

which results in muscle control functions that are essentially random and achieve very poor locomotion. Nearer sharks have learned for progressively longer periods of time (1350 more annealing steps). After learning for 6840 annealing steps, the nearest shark locomotes very efficiently and it wins the race decisively.

4.3 Abstraction of High-Level Controllers

In the second stage of the learning process abstracts higher-level muscle controllers. This is a dimensionality reducing change of representation that compresses the information content of the many control points to a compact form in terms of a few global basis functions. Specifically, it tries to represent the control functions as accurately as possible in the form $\mathbf{u}(t) \approx \sum_{i=0}^M \alpha_i \mathbf{B}_i(t)$, where \mathbf{B}_i are basis functions, α_i are scalar quantities, and M is a small number.

Since natural locomotion patterns are generally periodic [18], the Fourier basis is a reasonable choice. We employ the short time FFT to perform the change in basis (the wavelet transform may also be applied). If the Fourier space is a suitable representation that captures the temporal structure of the control functions, the dimensionality reduction can be achieved easily by eliminating all basis functions whose coefficients α_i in the above approximation formula are negligible. This will result in a small set of M significant basis functions, usually 1 or 2, with associated coefficients that define the abstracted muscle controller.

The artificial fish can now train itself to perform several routine locomotion tasks, such as swimming forward at different speeds and executing turns of different radii. Fig. 8 illustrates the control abstraction procedure and its results after straight swim and left turn training sessions. After it has abstracted controllers for these tasks, the artificial fish can construct a forward speed and steering map by interpolating across amplitudes, frequencies, and phases of the sinusoidal basis functions in the set of abstracted controllers. Finally, it can put these learned abstractions to use to accomplish higher level tasks, such as target tracking. Fig. 9 shows a shark model that has trained itself to swim between targets (spherical buoys); having swum from the left to the far target, it has now turned and is proceeding to the near target.




Figure 7: Race between sharks that have learned to locomote for progressively longer times.

5 Sensory Perception

The perception system of the artificial fish, illustrated in Fig. 1, comprises a set of virtual on-board sensors and a perceptual focuser. Currently the artificial fish is equipped with two sensors that provide information about the dynamic environment—a temperature sensor that measures the ambient (virtual) water temperature at the center of the artificial fish’s body and a much more elaborate cyclopean vision sensor.

5.1 Vision Sensor

We have not attempted to emulate the vision system of real fishes. Instead, we have incorporated a basic cyclopean vision sensor into the fish model. It is crucial to model the basic limitations of animal vision systems, otherwise the perceptually driven behaviors will not be natural. The cyclopean vision sensor has a 300 degree spherical field of view extending frontally and laterally to an effective radius V_r appropriate to the visibility of the translucent water (Fig. 10(a)). An object is “seen” only if some part of it enters this view volume and it is not fully occluded behind some other opaque object (Fig. 10(b)).

The artificial fish’s vision sensor has access to the geometry, material property, and illumination information that is available to the graphics pipeline for rendering purposes. In addition, the vision sensor can interrogate the world model database to identify nearby objects and interrogate the physical simulation to obtain information such as the instantaneous positions and velocities of objects of interest. In this way, the vision sensor extracts from the 3D virtual world only some of the most useful information that piscine visual processes can provide real fishes about their world, such as overall brightness, and the colors, sizes, distances, and identities of visible objects.

A more biologically plausible emulation of piscine visual processes would involve the application of various computer vision algorithms [10] to extract information from “retinal” images of the 3D world rendered from the vantage point of the artificial fish’s cyclopean vision sensor. Intrinsic images are a useful paradigm for this purpose [4]. The rendering pipelines of 3D graphics workstations can readily synthesize retinal images, associated z-buffers, and object identity maps for artificial fish vision. Currently, the fish determines only the overall brightness of its environment by computing the mean intensity of the retinal image. Fig. 11 shows examples of retinal images acquired by a fish “witnessing” another fish being baited by a fishing line.

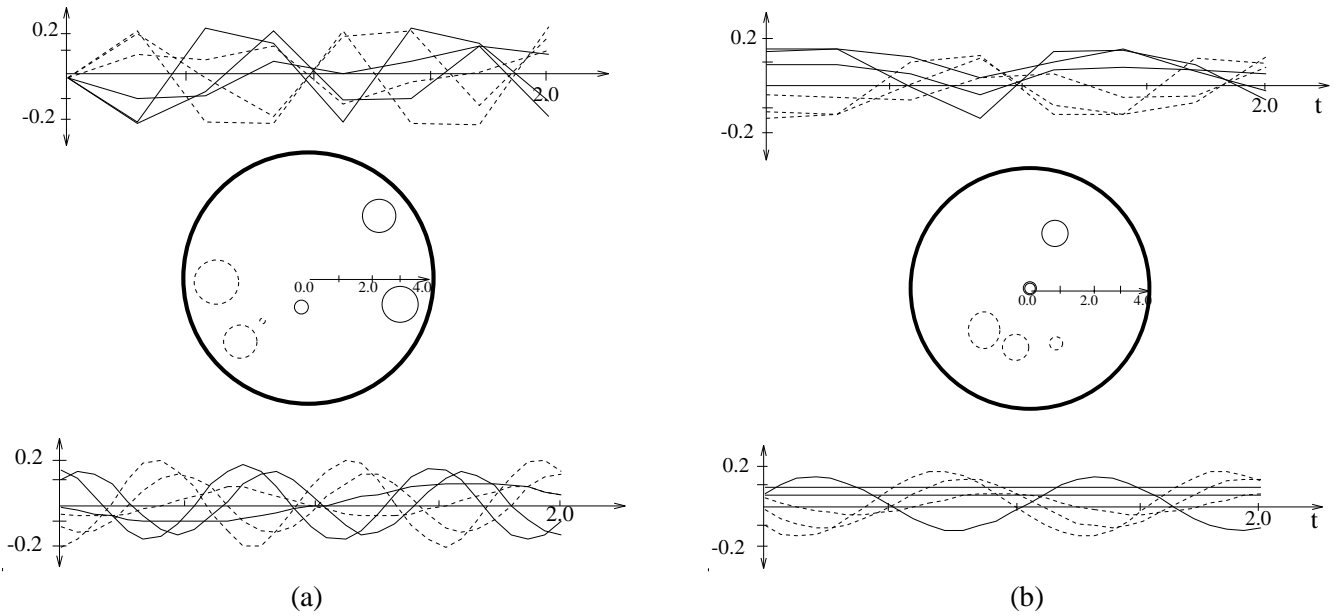


Figure 8: Abstraction of muscle controllers in straight swim (a) and left turn (b) training sessions. Upper plots: muscle controls gleaned by low-level learning algorithm. Center plots: primary modes. Lower plots: compact muscle controls abstracted from primary modes.

6 Behavioral Modeling

The artificial fish's behavior system runs continuously within the simulation loop. At each time step the intention generator issues an intention based on the fish's habits, mental state, and incoming sensory information. It then chooses and executes a behavior routine which in turn runs the appropriate motor controllers. It is important to note that the behavior routines are incremental by design. Their job is to bring the artificial fish one step closer to fulfilling the intention during the current time step. The intention generator employs a memory mechanism to avoid dithering.

6.1 Habits and Mental State

The innate character of the fish is determined by a set of habit parameters that determine whether or not it likes brightness, darkness, cold, warmth, schooling, or is a male/female, etc.

The artificial fish has three mental state variables, hunger H , libido L , and fear F . The range of each variable is $[0, 1]$, with higher values indicating a stronger urge to eat, mate and avoid danger, respectively. The variables are calculated as follows:

$$\begin{aligned}
 H(t) &= \min[1 - n^e(t)R(\Delta t^H)/\alpha, 1], \\
 L(t) &= \min[s(\Delta t^L)(1 - H(t)), 1], \\
 F(t) &= \min\left[\sum_i F^i, 1\right], \text{ where } F^i = \min[D_0/d^i(t), 1];
 \end{aligned}$$

where t is time, $n^e(t)$ is the amount of food consumed as measured by the number of food particles or prey fishes eaten, $R(x) = 1 - p_0x$ with constant p_0 is the digestion rate, Δt^H is the time since the last meal, α is a constant that dictates the appetite of the fish (bigger fishes have a larger α), $s(x) = p_1x$ with constant p_1 is the libido function, Δt^L is the time since the last mating, $D_0 = 100$ is a constant, and F^i and d^i are, respectively, the fear of and distance to sighted predator i . Nominal constants are $p_0 = 0.00067$ and $p_1 = 0.0025$. Certain choices can result in ravenous fishes (e.g, $p_0 = 0.005$) or sexual mania (e.g., $p_1 = 0.01$).



Figure 9: Trained shark swimming between targets.

6.2 Intention Generator

Fig. 12 illustrates the generic intention generator which is responsible for the goal-directed behavior of the artificial fish in its dynamic world.

The intention generator first checks the sensory information stream to see if there is any immediate danger of collision. If any object penetrates the fish's collision sensitivity region (a bounding box) then the intention I generated is to *avoid* collision. A large sensitivity region results in a 'timid' fish that takes evasive action to avoid a potential collision well in advance, while a tight sensitivity region yields a 'courageous' fish that takes evasive action only at the last second.

If there is no immediate danger of collision, the neighborhood is searched for predators, the fear state variable F and the most dangerous predator m for which $F^m \geq F^i$ are calculated. If the total fear $F > f_0$ (where $0.1 \leq f_0 \leq 0.5$ is a threshold value) evasive action is to be taken. If the most dangerous predator is not too threatening (i.e. $F^m < f_1$ where $f_1 > f_0$) and the fish has a schooling habit, then the *school* intention is generated, otherwise the *escape* intention is generated.

If fear is below threshold, the hunger and libido mental state variables H and L are calculated. If the greater of the two exceeds a threshold $0 < r < 0.5$, the intention generated will be to *eat* or *mate* accordingly.

If the above test fails, the intention generator accesses the ambient light and temperature information from the perception system. If the fish's habits dictate contentment with the ambient conditions, the intention generated will be to *wander* about, otherwise it will be to *leave* the vicinity.

Note that after the intention generator chooses an intention, it invokes the perceptual focus mechanism. For example, when the *avoid* intention is generated, the perception focuser is activated to locate the positions of the obstacles, paying special attention to the most dangerous one, generally the closest. Then the intention generator computes motor preferences (qualitative constraints, such as *another fish to the left* \Rightarrow *no left turn*). The focuser passes only the position of the most dangerous obstacle along with these constraints to the behavior routines. When the intention of a male fish is to *mate*, the focuser targets the most desirable female fish; when the intention is to *escape* from predators, only the information about the most threatening predator is passed to the next layer; etc.

6.3 Behavior Memory and Persistence

In a complex dynamic world, the artificial fish should have some persistence in its intentions, otherwise it will tend to dither, perpetually switching goals. If the current behavior is interrupted by a high priority event, the intention generator stores in a single-item short term memory the current intention and some associated information that may




Figure 10: Artificial fish vision sensor. (a) Visual perception is limited to 300 degree solid angle. (b) Occlusion and distance limits the perception of objects (only the fish towards the left is visible).

be used to resume the interrupted behavior. Persistence is particularly important in making long duration behaviors such as foraging, schooling, and mating more robust. Suppose, for example, that the current behavior is mating and an imminent collision is detected with another fish. This causes an *avoid* intention and the storage of the *mate* intention (we refer to the stored intention as I^s) along with the identity of the mating partner. After the obstacle has been cleared, the intention generator commands the focuser to generate up-to-date heading and range information about the mating partner, assuming it is still in viewing range.

Our design of the intention generator and focuser simplifies the modification of existing personalities and behaviors and the addition of new ones. For example, we can create artificial fishes with different persistences by augmenting the focuser with a new positive threshold. Suppose the current intention of a predator fish is to *eat* and let the distance to some currently targeted prey be l_c and the distance to some other prey be l_n . If $l_c - l_n$ is greater than the threshold, the fish will target the new prey. Varying the threshold will vary the fish's level of persistence. The same heuristic can be applied to mates when the fish is trying to *mate*. One can make the fish 'fickle' by setting the value of the threshold close to zero or make it 'devoted' by setting a large value.

6.4 Behavior Routines

Once the intention generator selects an intention it attempts to satisfy the intention by passing control to a behavior routine along with the data from the perception focuser. The artificial fish currently includes eight behavior routines: *avoiding-static-obstacle*, *avoiding-fish*, *eating-food*, *mating*, *leaving*, *wandering*, *escaping*, and *schooling* which serve the obvious purposes. The behavior routine uses the focused perceptual data to select an MC and provide it with the proper motor control parameters. We now briefly describe the function of the routines.

The *avoiding-static-obstacle* and *avoiding-fish* routines operate in similar fashion. Given the relative position of the obstacle, an appropriate MC (e.g. *left-turn-MC*) is chosen and the proper control parameters are calculated subject to the motor preferences imposed by other surrounding obstacles. For efficiency the *avoid-fish* routine treats the dynamic obstacle as a rectangular bounding box moving in a certain direction. Although collisions between fishes cannot always be avoided, bounding boxes can be easily adjusted such that they almost always are, and the method is very efficient. An enhancement would be to add collision resolution.

The *eating-food* routine tests the distance d from the fish's mouth to the food (see Fig. 5). If d is greater than some

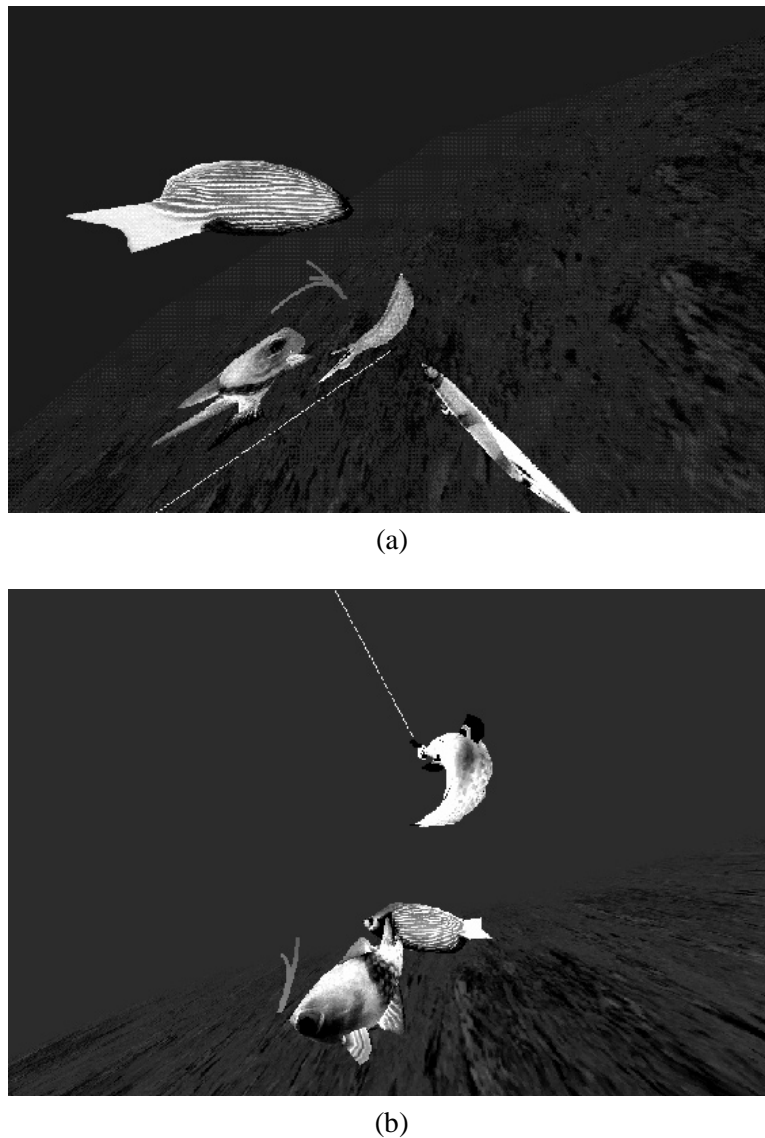


Figure 11: Fisheye view of the world showing fishing line (a) and hooked fish (b).

threshold value, the subroutine *chasing-target* is invoked.³ When d is less than the threshold value the subroutine *suck-in* is activated where a “vacuum” force (to be explained in Sec. 6.1) is calculated and then exerted on the food.

The *mating* routine invokes four subroutines: *looping*, *circling*, *ascending* and *nuzzling* (see Sec. 7.3 for details). The *wandering* routine sets the fish swimming at a certain speed by invoking the swim-MC, while sending random turn angles to the turn-MCs. The *leaving* routine is similar to the *wandering* routine. The *escaping* routine chooses a suitable MC according to the relative position, orientation of the predator to the fish. The *schooling* routine will be discussed in Sec. 7.2.

7 Artificial Fish Types

In the introductory paragraph of the paper we described the behavior of three types of artificial fishes—predators, prey, and pacifists. This section presents their implementation details.

³The *chasing-target* subroutine guides a fish as it swims towards a goal. It plays a crucial role in several behavior routines.

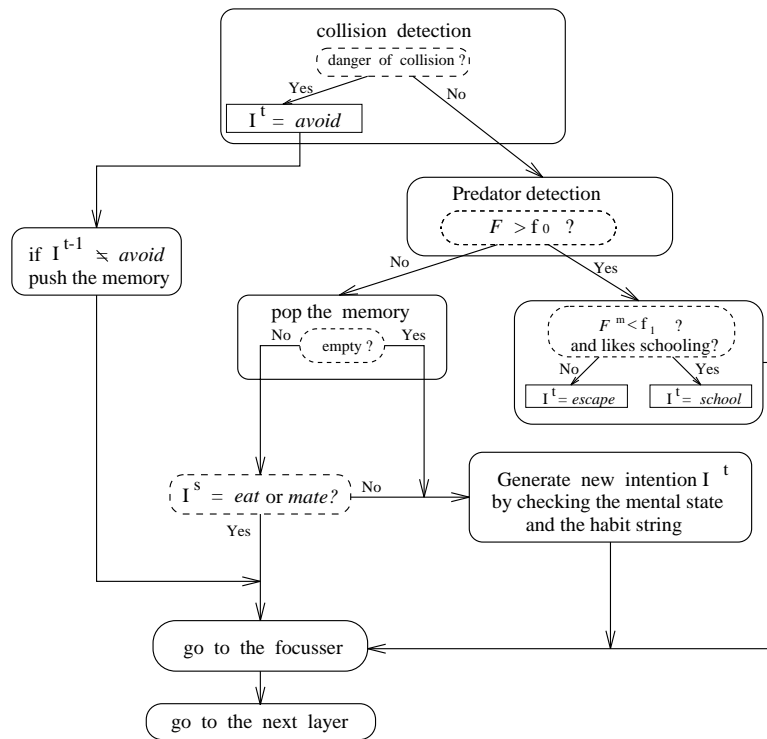


Figure 12: Generic intention generator (simplified). Set of intentions: $\{ avoid, escape, school, eat, mate, leave, wander \}$. f_0 and f_1 are thresholds with $f_0 < f_1$.

7.1 Predators

Fig. 13 is a schematic of the intention generator of a predator, which is a specialized version of Fig. 12. For simplicity, predators currently are not preyed upon by other predators, so they perform no predator detection, and *escape*, *school*, and *mate* intentions are disabled ($F = 0$, $L = 0$). Since predators cruise perpetually, the *leave* intention is also disabled.

Generally prey is in less danger of being hunted when it is far away from the predator, or is in a school, or is behind the predator. A predator chases prey k if the cost $C_k = d_k(1 + \beta_1 S_k + \beta_2 E_k/\pi)$ of reaching it is minimal. Here, d_k is the distance between the mouth of the predator and the center of prey k 's body, $S_k = 1$ if prey k is in a school of fishes, otherwise $S_k = 0$, and the angle $E_k \in [0, \pi)$ (Fig. 5) measures the turning cost. β_1 and β_2 are parameters that tune the contributions of S_k and E_k . We use $\beta_1 = 0.5$ and $\beta_2 = 0.2$ in our implementation of the focuser. Plate 1c shows a shark predator stalking a school of prey fish.

Most teleost fishes do not bite on their victims like sharks do. When a fish is about to eat it swims close to the victim and extends its protrusile jaw, thus creating a hollow space within the mouth. The pressure difference between the inside and the outside of the mouth produces a vacuum force that sucks the victim and anything else in the nearby water into the mouth. The predator closes its mouth, expels the water through the gills, and grinds the food with pharyngeal jaws [25]. We simulate this process by enabling the artificial fish to open and close its mouth kinematically. To suck in prey, it opens its mouth and, while the mouth is open, exerts vacuum forces on fishes (the forces are added to external nodal forces \mathbf{f}_i in equation (1) and other dynamic particles in the vicinity of the open mouth, drawing them in (Fig. 14).

7.2 Prey

The intention generator of a prey fish is given by specializing the generic intention generator of Fig. 12, as shown in Fig. 15.

Schooling and evading predators are the two distinct behaviors of prey. We briefly describe the implementation of the *schooling* behavior. Schooling is a complex behavior where all the fishes swim generally the same direction.

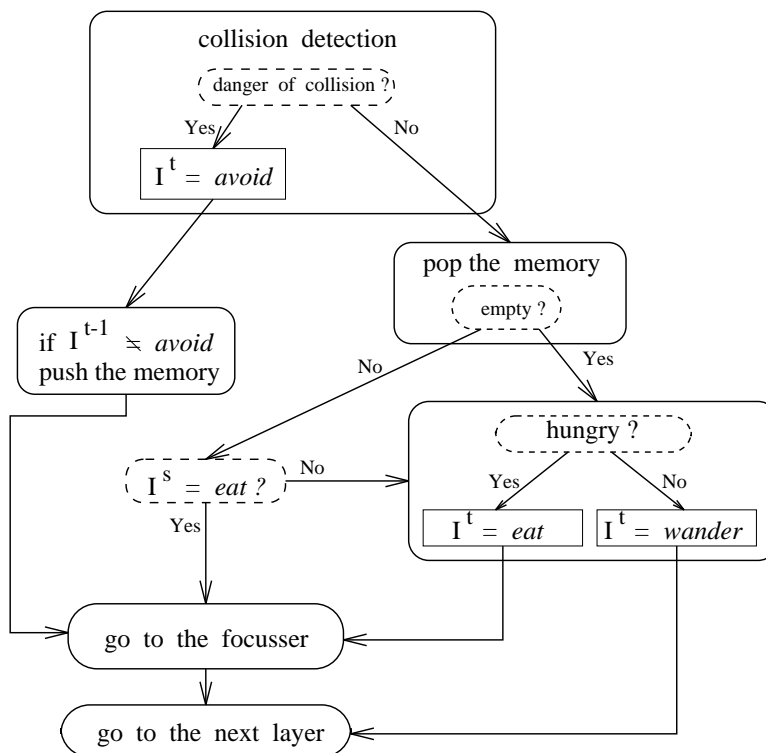


Figure 13: The intention generator of a Predator

Each fish constantly adjusts its speed and direction to match those of other members of the school. They establish a certain distance from one another, roughly one body length from neighbors, on average [25]. Each member of a school of artificial fish acts autonomously, and the schooling behavior is achieved through sensory perception and locomotion. An incentive school is formed when a few fish swim towards a lead fish. Once a fish is in some proximity to some other schooling fish, the *schooling* behavior routine outlined in Fig. 16 is invoked.

The intention generator prevents schooling fish from getting too close together, because the *avoid* collision intention has highest precedence. To create more compact schools, the collision sensitivity region of a schooling fish is decreased, once it gets into formation. When a large school encounters an obstacle, the autonomous behavior of individual fishes trying to avoid the obstacle may cause the school to split into two groups and rejoin once the obstacle is cleared and the *schooling* behavior routine regains control (Fig. 17).

7.3 Pacifists

The intention generator of a pacifist differs from that of prey in that intention *mate* is activated and *escape* and *school* are deactivated.

Piscine mating behaviors show great interspecies and intraspecies diversity [21]. However, two behaviors are prevalent: (i) nuzzling, where typically the male approaches the female from underneath and nudges her abdomen repeatedly until she is ready to spawn, and (ii) spawning ascent, where in its simplest form, the female rapidly swims towards the surface pursued by the male and releases gametes at the peak of her ascent. Moreover, courtship dancing is common in many species, albeit with substantial variation. Two frequently observed patterns are looping, in which the male swims vigorously up and down in a loop slightly above and in front of the female, and circling, in which the male and female circle, seemingly chasing each other's tail.

We have implemented a reasonably elaborate *courtship* behavior routine which simulates courtship dancing, circling, spawning ascent, and nuzzling behavior patterns in sequence (Plate 1b). A male fish selects a mating partner based on the following criteria: a female of the same species is more attractive than one of different species, and closer females are more attractive than ones further away. A female selects a partner similarly, but shows preference over the size of the male fish (stronger, more protective) rather than its distance.



Figure 14: Predator ingesting prey.




Figure 15: Modified portion of intention generator for prey.

Once fish i has selected a potential partner j based on the above criteria, it sends a signal to fish j , and there are three possibilities: *Case 1*: If fish j 's intention is not to *mate*, fish i approaches j and follows it around using *chasing-target* with the center of j 's body as target. *Case 2*: If fish j 's intention is to *mate* but its intended partner is not fish i . In this case, if i is male it will perform a *looping* behavior in front of j for a certain amount of time. If j is impressed and selects i during this time limit, then the courtship sequence continues, otherwise i will discontinue *looping* and leave j to find a new potential partner. Otherwise, if i is female it will choose another potential male. *Case 3*: If fish j 's intention is to *mate* and its intended partner is fish i , the *courtship* behavior starts with the male looping in front of the female while she hovers and bobs her head. Looping is simulated by invoking *chasing-target* at a point in front of the female's head which moves up and down at a certain frequency. The female's hovering and head bobbing is accomplished through motor control of her pectoral fins (i.e., parameter γ in Fig. 5).

The male counts the number of times his mouth reaches the vicinity of the moving point, and when the count exceeds a set threshold (currently 6) he makes a transition from *looping* to *circling* behavior. Although the threshold count is fixed, the actual motions and duration of looping is highly unpredictable for any number of reasons, including the fact that looping may be temporarily interrupted to handle high priority events such as potential collisions between the pair or with other fishes that may pass by.

Before the transition to *circling*, the female fish may reject her initial partner and turn to a new larger male fish if the latter joins in the *looping* display. At this point the initially engaged male turns away as in case 2 described above. *Circling* is achieved when the fishes use *chasing-target* to chase each other's tail.

The *circling* routine ends and the spawning *ascending* routine begins after the female has made a fixed number of turns during circling. The female fish ascends quickly through fast swimming followed by hovering. The male

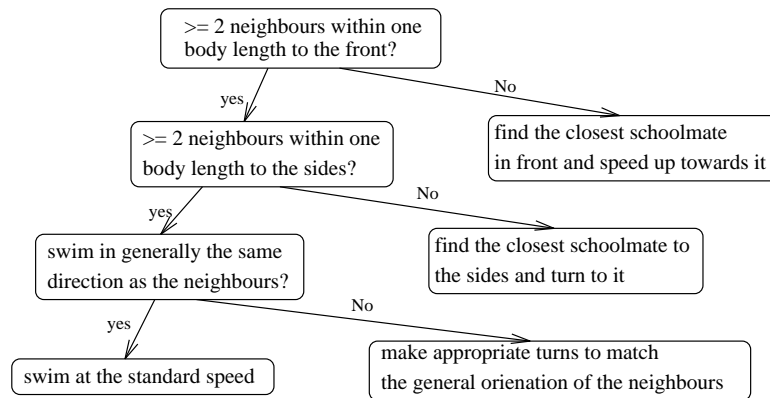


Figure 16: Schooling behavior routine.




Figure 17: School of fish swimming past cylindrical obstacle (from lower left to upper right). Schooling behavior is interrupted by collision avoidance behavior and then resumed.

fish uses *chasing-target* to follow the abdomen of the female. The *nuzzling* routine requires the male to approach her abdomen from below. Once his mouth touches her abdomen, the male backs off for a number of time steps. This procedure repeats, until the male successfully touches the female a predetermined number of times (e.g., 3). To permit the mating pair to come close together, the regions of sensitivity are set very tightly to their bodies. It is intriguing to watch some of the male artificial fish's nuzzling attempts fail because of an inappropriate approach angle to the female which triggers the *avoiding-fish* response. The male turns away to avoid the collision and tries again.

8 Conclusion and Research Directions

This paper has presented the results of research spanning the fields of artificial life and computer graphics, with anticipated future implications for computer vision. We have developed a physics-based, virtual marine world inhabited by astonishingly lifelike artificial life forms that emulate the appearance, motion, and behavior of fishes in their natural habitats. Each artificial fish is an autonomous agent with a deformable body actuated by internal muscles, eyes, and a mind that includes learning, behavior, perception, and motor centers. Through controlled muscle

actions, artificial fishes are able to swim through simulated water in accordance with simplified hydrodynamics. Their functional fins enable them to locomote, maintain balance, and maneuver in the water. Though rudimentary compared to real animals, their minds are nonetheless able to learn some basic motor functions and carry out perceptually guided motor tasks. In accordance with their perceptual awareness of the virtual world, their minds arbitrate a repertoire of piscine behaviors, including collision avoidance, foraging, preying, schooling, and mating. The easy extensibility of our approach is suggested most evidently by the complex patterns of mating behavior that we have been able to implement in artificial fishes.

Our model achieves a good compromise between realism and computational efficiency. To give an example simulation rate, our implementation can simulate a scenario with 10 fishes, 15 food particles, and 5 static obstacles at about 4 frames/sec, including wireframe rendering time, on a Silicon Graphics R4400 Indigo² Extreme workstation. More complex scenarios with large schools of fish, dynamic plants, and full color texture mapped rendering at video resolution can take 5 seconds or more per frame.

Our work opens up many avenues of research. Clearly the artificial fish is a virtual robot that offers a much broader range of perceptual and animate capabilities, lower cost, and higher reliability than can be expected from present-day physical robots like those described in [13]. For at least these reasons, artificial fishes in their dynamic world can serve as a proving ground for theories that profess competence at effectively linking perception to action [3]. A different research direction would address the goals of researchers interested in evolving artificial life. We may be within reach of computational models that can imitate the spawning behaviors of the female fish and fertilization by the male. Through simulated sexual reproduction in a competitive world, gametes representing artificial fish genotypes can be fused to evolve new varieties of artificial fishes. Interestingly, Pokhilko, Pajitnov, *et al.*, have already demonstrated the simulated breeding of fish models much simpler than ours using genetic algorithms, and this idea has resulted in the fascinating computer game “El-Fish” [9].

Acknowledgements

We thank Eugene Fiume, Michael McCool, Michiel van de Panne, Sarah Peebles, and John Funge for discussions and for assistance with the computer animations. This research has been made possible by a grant from the Natural Sciences and Engineering Research Council of Canada and the support of the Canadian Institute for Advanced Research.

References

- [1] H. E. Adler. *Fish Behavior: Why Fishes do What They Do*. T.F.H Publications, Neptune City, NJ, 1975.
- [2] R.M. Alexander. *Exploring Biomechanics*. Scientific American Library, New York, 1992.
- [3] D. Ballard. Animate vision. *Artificial Intelligence*, 48:57–86, 1991.
- [4] H. G. Barrow and J. M. Tenenbaum. Recovering intrinsic scene characteristics from images. In E. Riseman and A. Hanson, editors, *Computer Vision Systems*, pages 3–26. Academic Press, NY, 1978.
- [5] R. Beer. *Intelligence as Adaptive Behavior*. Academic press, NY, 1990.
- [6] R. W. Blake. *Fish Locomotion*. Cambridge University Press, Cambridge, England, 1983.
- [7] V. Braitenberg. *Vehicles, Experiments in Synthetic Psychology*. MIT Press, Cambridge, MA, 1984.
- [8] R. A. Brooks. Elephants don’t play chess. *Robotics and Autonomous Systems*, 6(1):3–15, 1990.
- [9] E. Corcoran. One fish, two fish: How to raise a school of tempting software toys. *Scientific American*, July 1992.
- [10] B. K. P. Horn. *Robot Vision*. MIT Press, Cambridge, MA, 1986.

- [11] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int. J. of Computer Vision*, 1(4):321–331, 1987.
- [12] K. Lorenz. *Foundations of Ethology*. Springer-Verlag, New York, 1973.
- [13] P. Maes, editor. *Designing Autonomous Agents*. MIT Press, Cambridge, MA, 1991.
- [14] M. J. Matarić. *Interaction and Intelligent Behavior*. PhD thesis, Dept. of EECS, MIT, Cambridge, MA, May 1994.
- [15] D. McFarland. *Animal Behaviour*. Pitman, 1985.
- [16] J.-A. Meyer and A. Guillot. Simulation of adaptive behavior in animats: Review and prospect. In J.-A. Meyer and S. Wilson, editors, *From Animals to Animats*, pages 2–14. MIT Press, Cambridge, MA, 1991.
- [17] G. S. P. Miller. The motion dynamics of snakes and worms. *Computer Graphics*, 22(4):169–177, 1988.
- [18] K. G. Pearson. Sensory elements in pattern-generating networks. In *Making Them Move*, pages 111–127. Morgan Kaufmann, San Mateo, California, 1991.
- [19] W. Press, B. Flannery, S. Teukolsky, and W. Vetterling. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, Cambridge, England, 1986.
- [20] C. W. Reynolds. Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21(4):25–34, 1987.
- [21] R. E. Thresher. *Reproduction in Reef Fishes*. T.F.H. Publications, Neptune City, NJ, 1984.
- [22] N. Tinbergen. *The Study of Instinct*. Clarendon Press, Oxford, England, 1950.
- [23] X. Tu, D. Terzopoulos, and E. Fiume. Go Fish! ACM SIGGRAPH Video Review Issue 91: SIGGRAPH'93 Electronic Theater, 1993.
- [24] P. W. Webb. Form and function in fish swimming. *Scientific American*, 251(1), 1989.
- [25] R. Wilson and J. Q. Wilson. *Watching Fishes*. Harper and Row, New York, 1985.
- [26] S. W. Wilson. The animat path to AI. In J.-A. Meyer and S. Wilson, editors, *From Animals to Animats*, pages 15–21. MIT Press, Cambridge, MA, 1991.