

Artificial Fishes: Physics, Locomotion, Perception, Behavior

Xiaoyuan Tu and Demetri Terzopoulos
Department of Computer Science, University of Toronto¹

Keywords: *behavioral animation, artificial life, autonomous agents, animate vision, locomotion control, physics-based modeling*

Abstract: *This paper proposes a framework for animation that can achieve the intricacy of motion evident in certain natural ecosystems with minimal input from the animator. The realistic appearance, movement, and behavior of individual animals, as well as the patterns of behavior evident in groups of animals fall within the scope of the framework. Our approach to emulating this level of natural complexity is to model each animal holistically as an autonomous agent situated in its physical world. To demonstrate the approach, we develop a physics-based, virtual marine world. The world is inhabited by artificial fishes that can swim hydrodynamically in simulated water through the motor control of internal muscles that motivate fins. Their repertoire of behaviors relies on their perception of the dynamic environment. As in nature, the detailed motions of artificial fishes in their virtual habitat are not entirely predictable because they are not scripted.*

1 Introduction

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 gracefully swim 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 hunt for smaller prey fishes. Prey fishes swim around happily until they see a predator, at which point they take evasive action. When a predator appears in the distance, similar species of prey form schools to improve their chances of escape. When a predator approaches 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 forage on floating

plankton when they are hungry. When compelled by their libidos, they engage in elaborate courtship rituals to secure mates.

The animation of such scenarios with visually convincing results has been elusive. In this paper, we develop an animation framework within whose scope fall all of the above complex patterns of action, and many more, without any keyframing. The key to achieving this level of complexity, and beyond, with minimal intervention by the animator, is to create fully functional artificial animals—in this instance, artificial fishes. Artificial fishes are autonomous agents whose appearance and complicated group interactions are as faithful as possible to nature’s own. To this end, we pursue a bottom-up, compositional approach in which we model not just form and superficial appearance, but also the basic physics of the animal and its environment, its means of locomotion, its perception of its world, and last but not least, its behavior. The holistic nature of our approach to synthesizing artificial fishes is crucial to achieving realism. Partial solutions that do not adequately model physics, locomotion, perception, and behavior, and do not combine these models intimately within the agent will not produce convincing results.

An early result of our research is the computer animation “Go Fish!” [18]. The final sequence of this animation shows a colorful variety of artificial fishes feeding in translucent water (see Plate 1). Dynamic aquatic plants grow from the seabed. A sharp hook on a line descends towards the hungry fishes and attracts them (Plate 1(a)). A hapless fish, the first to bite the bait, is caught and dragged to the surface (Plate 1(b)). Only the camera and the fishing line were scripted in the animation. The beauty of the animation is enhanced by the detailed motions of the artificial fishes which emulate the complexity and unpredictability of movement of their natural counterparts.

1.1 Background

At its lowest level of abstraction, our work is an instance of physics-based graphics modeling. The physics-based artificial fish model that we develop is inspired by the surprisingly effective model of snake and worm dynamics proposed by Miller [9] and the facial model proposed by Terzopoulos and Waters [14]. Our fish model is also an animate spring-mass system with internal contractile muscles that are activated to produce the desired motions. Unlike these previous models, however, we simulate the spring-mass system using a more sophisticated implicit Euler method which maintains the stability of the simulation over the large dynamic range of forces produced in our simulated aquatic world. Using spring-mass systems, we also model the dynamic plants found in the artificial fish habitat.

The control of physics-based animate models has attracted significant attention, especially the control of articulated models to animate legged locomotion [2, 11]. Fish animation poses control challenges particular to highly deformable, muscular bodies not un-

¹10 King’s College Road, Toronto, Ontario, Canada, M5S 1A4
E-mail: {tu|dt}@cs.toronto.edu

like those of snakes [9]. We have devised a motor control subsystem that achieves muscle-based, hydrodynamic locomotion by simulating the forces of interaction of a deformable body in an aquatic medium. The motor controller harnesses the hydrodynamic forces produced by fins to achieve forward locomotion over a range of speeds, execute turns, and alter body roll, pitch, and yaw so that the fish can move at will within its 3D virtual world.

At a higher level of abstraction, our research may be categorized as advanced behavioral animation. Several researchers have endeavored to develop behavioral models for computer animation. Simpler behavioral approaches than ours have been proposed to cope with the complexity of animating anthropomorphic figures [22], animating the synchronized motions of flocks, schools, or herds [13], and interactive animation control [20] (see, also, the papers by Badler, Calvert, Girard, Green, Miller, Wilhelms, and Zeltzer in [2]). Artificial fishes are “self-animating” in the sense of Reynolds’ pioneering work [13], but unlike his procedural “boi” actors, they are fairly elaborate physical models.

To achieve a level of behavioral realism consistent with the locomotive abilities of artificial fishes, it is prudent to consult the ethology literature [16, 5, 7, 1]. Tinbergen’s landmark studies of the three-spined stickleback highlight the great diversity of piscatorial behavior, even within a single species. We achieve the nontrivial patterns of behavior outlined in the introductory paragraph of this paper, 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 [3]. Then we combine the primitive behaviors into motivational behaviors whose activation depends also on the artificial fish’s mental state, including hunger, libido, etc.

Behavior is supported by perception as much as it is by action. Evolution has developed in most animals, including fishes, acute perceptual modalities to increase their chances of survival in an unpredictable and often hostile world. Reynolds’ “boi” maintained flock formations through simple perception of other nearby actors [13]. The roach actor described in [8] retreated when it sensed danger from a virtual hand. Renault *et al.* [12] advocate a more extensive form of synthetic vision for behavioral actors, including the automatic computation of internal spatial maps of the world. 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. Furthermore, they are equipped with secondary nonvisual modalities, such as the ability to sense the local water temperature.

The confluence of behavior, perception, and motor systems makes the artificial fish an autonomous agent. In this regard, our design of virtual agents is compatible with recent work in robotics aimed at the implementation of physical agents (see, e.g., the compilation [6]). Interestingly, as our holistic computational model exceeds a certain level of physical, motor, perceptual, and behavioral sophistication, the agent’s range of functionality broadens due to emergent behaviors, not explicitly programmed, but nonetheless manifest as the artificial fish interacts with a complex dynamic world populated by other artificial fishes. We aim to emulate convincingly the appearance of the animal as well, so that our computational model will be useful for the purposes of animation.

1.2 Overview

Fig. 1 shows an overview of an artificial fish situated in its world, illustrating the motor, perception, and behavior subsystems.

The motor system comprises the dynamic model of the fish, the actuators, and a set of motor controllers (MCs). Since our goal is to animate an animal realistically and at reasonable computational cost, we have crafted a mechanical model that represents a good compromise between anatomical consistency, hence realism, and

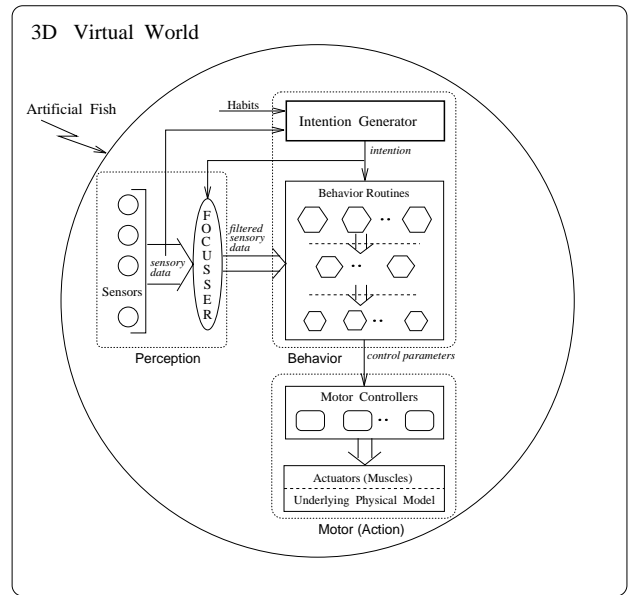


Figure 1: Control and information flow in artificial fish.

computational efficiency. Our model is rich enough so that we can build MCs by gleaning information from the animal biomechanics literature. 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 employs a set of virtual on-board sensors to provide all the sensory information about the dynamic environment. The system includes a perceptual attention mechanism which allows the artificial fish to train its sensors at the world in a task-specific way, filtering out undesired sensory information according to the needs of the behavior routines.

The behavior system of the artificial fish mediates between its perception system and its motor system. An intention generator, the fish’s “cognitive” center, harnesses the dynamics of the perception-action cycle. The animator establishes the innate character of the fish through a set of habit parameters that determine whether or not it likes darkness or is a male/female, etc. The intention generator combines the habits 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 to filter out sensory information unnecessary to accomplishing the goal at hand. For example, if the intention is to eat food, then the artificial fish attends to sensory information related to nearby food sources. At every simulation time step, the intention generator activates behavior routines that input the filtered sensory information and compute the appropriate motor control parameters to carry the fish one step closer to fulfilling the current intention. Primitive behavior routines, such as obstacle avoidance, and more sophisticated motivational behavior routines, such as mating, implement the artificial fish’s repertoire of behaviors.

2 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 [19]. 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. 2).




Figure 2: The spring-mass dynamic fish model. Springs are at their rest lengths.

2.1 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 spring S_{ij} connect node i to node j . 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 $\mathbf{r}_{ij} = \mathbf{x}_j(t) - \mathbf{x}_i(t)$. The force S_{ij} exerts on node i is $\mathbf{f}_{ij}^s = c_{ij}e_{ij}(t)\mathbf{r}_{ij}/\|\mathbf{r}_{ij}\|$ (and it exerts the force $-\mathbf{f}_{ij}^s$ on node j). The Lagrange equations of motion of the dynamic fish are:

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

where ϱ_i is the damping factor, $\mathbf{w}_i(t) = \sum_{j \in N_i} \mathbf{f}_{ij}^s(t)$ 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 is the external (hydrodynamic) force on node i .

To integrate the differential equations of motion, we employ a numerically stable, implicit Euler method [10]. The method assembles the sparse stiffness matrix for the spring-mass system in “skyline” storage format. The matrix is factorized once at the start of the simulation and then resolved at each time step.¹

2.2 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 characteristic swinging 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 swings, it sets in motion a volume of water. The inertia of the displaced water produces a reaction force normal to

¹In 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 $\varrho_i = 0.05$ in (1) and the time step used in the Euler time-integration procedure is 0.055.

the fish’s body and proportional to the volume of water displaced per unit time, which propels the fish forward (Fig. 3(a)). 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 fish model between the nodes and approximate the force on each planar triangle as $\mathbf{f} = \min[0, -A(\mathbf{n} \cdot \mathbf{v})\mathbf{n}]$, 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 $\mathbf{f}/3$.

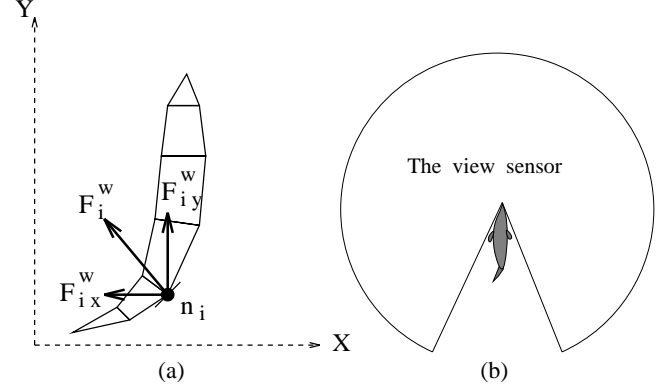


Figure 3: Hydrodynamic locomotion and vision sensor. (a) With tail swinging towards positive X axis, reaction force \mathbf{F}_i^w at point n_i acts along the inward normal. Component \mathbf{F}_{ix}^w resists the lateral movement, while \mathbf{F}_{iy}^w is forward thrust. Aggregate thrust propels fish towards positive Y axis. (b) Visual perception limited to 300 degree solid angle.

2.3 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 swimming MC controls the swimming segment muscles (see Fig. 2), while the turning MCs control the turning segment muscles.

According to [19], 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 suffice to control each of the two swim segments because of 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 models displays a sinusoidal body shape as the fish swims (see [19]).

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 [19]. The artificial fish turns by contracting and expanding the springs of the turning segments (Fig. 2) 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 effectively 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 (turn-MC(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.

2.4 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 serve as a brake by increasing drag [21]. Many reef fishes use a pectoral swimming style to achieve very fine motion control when foraging, including backwards motions, by keeping their bodies still and using their pectorals like oars.

The artificial fish 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. Instead of creating a detailed physics-based model of the pectoral fins, we are content to simulate only their dynamic effect on the locomotion of the fish. This is because for our purposes the detailed movement of the pectoral fins is of lesser interest than the movement of the fish body. Furthermore, we wish to simplify the fish model and its numerical solution.

The pectoral fins (Fig. 4) work by applying reaction forces to nodes in the midsection, i.e. nodes $1 \leq i \leq 12$ (see Fig. 2).




Figure 4: The pectoral fins

The pectoral fins are analogous to the wings 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. 4), the fin force is $F_f = -A(\|\mathbf{v}\| \cos \gamma) \mathbf{n} = -A(\|\mathbf{v}\| \cos \gamma) \mathbf{n}$ 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, for instance, in maintaining schooling patterns.

3 Sensory Perception

The artificial fish currently has two on-board sensors with which to perceive its environment and govern its actions—a vision sensor and a temperature sensor.

The temperature sensor samples the ambient water temperature at the center of the artificial fish’s body. The vision sensor is more complicated. We do not attempt to emulate the highly evolved vision system of a real fish. Instead, the vision sensor extracts from the 3D virtual world only some of the most useful information that fish vision can provide real fishes about their world, such as the colors, sizes, distances, and identities of objects.

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 object database to identify nearby objects and interrogate the physical simulation to obtain information such as the instantaneous velocities of objects of interest.

Currently, the artificial fish’s vision is cyclopean, and it covers a 300 degree spherical angle extending to an effective radius that is appropriate for the assigned visibility of the translucent water (Fig. 3(b)). An object is “seen” if any part of it enters this view volume and is not fully occluded by another object.

A more realistic emulation of piscatorial visual processes would involve the application of computer vision algorithms to extract information from images (and associated z-buffers) of the 3D world rendered from the vantage point of the artificial fish’s (binocular) vision sensor. At present the artificial fish can average the image to determine the overall light.

For further details about perceptual modeling in artificial fishes see [17].

4 Behavioral Modeling and Animation

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 get 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.

4.1 Habits

Using a simple user interface, the animator establishes the innate character of the fish through a set of habit parameters that determine whether it likes brightness, darkness, cold, warmth, schooling, is male or female, etc.

4.2 Mental State

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), $s(x) = p_1 x$ 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., $p_1 = 0.01$).

5 Intention Generator

Fig. 5 illustrates the generic intention generator which is responsible for the goal-directed behavior of the artificial fish in its dynamic world.

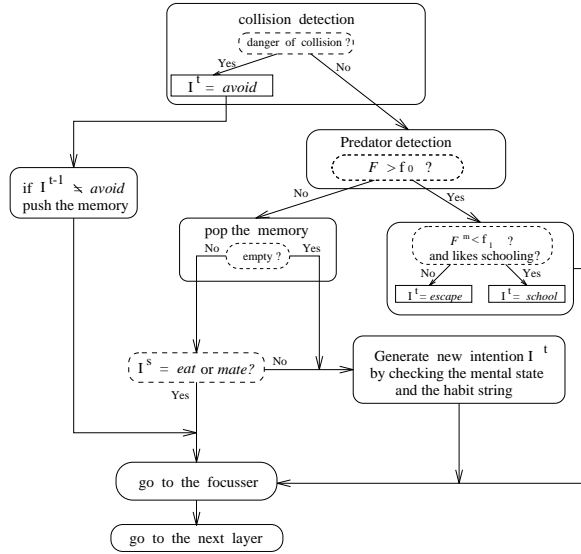


Figure 5: Generic intention generator (simplified). Set of intentions: $\{ avoid, escape, school, eat, mate, leave, wander \}$. f_0 and f_1 are thresholds with $f_0 < f_1$.

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 around the fish body) 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 focuser computes qualitative constraints, such as *obstacle 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.

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 is able to store, in a single-item short term memory, the current intention and some associated information that may be used to resume the interrupted behavior. Such persistence serves primarily to make longer duration behaviors such as feeding and mating more robust. Suppose for example that the current behavior is mating and an imminent collision is detected. 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 is 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. A similar scenario may occur during feeding.

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.

5.1 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 constraints 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. 4). If d is greater than some threshold value, the subroutine *chasing-target* is invoked.² When d is less than the

²The *chasing-target* subroutine guides a fish as it swims towards a goal. It plays a crucial role in several behavior routines. We describe it in more detail elsewhere [17].

threshold value the subroutine *suck-in* is activated where a “vacuum” force (to be explained in section 6.1) is calculated and then exerted on the food.

The *mating* routine invokes four subroutines: *looping*, *circling*, *ascending* and *nuzzling* (see Section 6.3 for details). The *wandering-about* 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-about* 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 Section 6.2.

6 Artificial Fish Types

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

6.1 Predators

Fig. 6 is a schematic of the intention generator of a predator, which is a specialized version of Fig. 5. To simplify matters, predators currently do not prey 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.

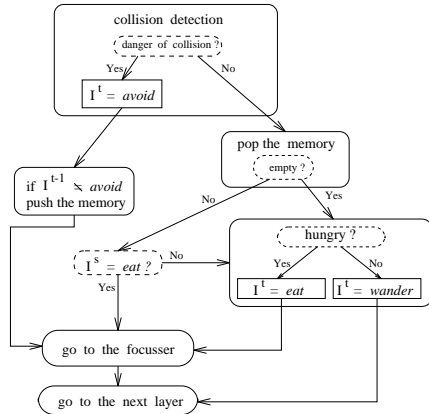


Figure 6: The intention generator of a Predator

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. 4) 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 2(a) shows a (green) predator stalking prey.

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 into the mouth the victim and anything else in the nearby water. The predator closes its mouth, expels the water through the gills, and grinds the food with pharyngeal jaws [21]. 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 f_i in

equation (1) and other dynamic particles in the vicinity of the open mouth, drawing them in (see Plate 2(b)).

6.2 Prey

The intention generator of a prey fish is a specialization of the generic intention generator of Fig. 5 as follows:




Figure 7: Portion of intention generator of prey.

Schooling and evading predators are two distinct behaviors of prey. Because of space constraints, we briefly describe the implementation of only the *schooling* behavior. Schooling is a complex behavior where all the fishes swim in generally the same direction (see Plate 3(a)). 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 [21]. As in [13], each member of a school of artificial fish acts autonomously, and the schooling behavior is achieved through sensory perception and locomotion. An inceptive school forms 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. 8 is invoked.

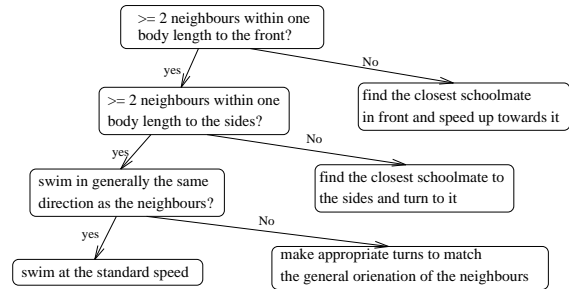


Figure 8: Schooling behavior routine.

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 obstacle avoidance behavior of individual fishes may cause the school to split into two groups and rejoin once the obstacle is cleared and the *schooling* behavior routine regains control (Plate 3(b)).

6.3 Pacifists

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

Piscatorial mating behaviors show great interspecies and intraspecies diversity [15]. 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 chase each other's tail.

We have implemented an elaborate *mating* behavior routine which simulates courtship competition (Plate 4(a)), looping, circling, spawning ascent, and nuzzling (Plate 4(b)) behavior patterns in sequence. A male fish selects a mate 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 mate similarly, but shows preference to male fish size (stronger, more protective) rather than proximity.

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)

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 items 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 above. *Circling* is achieved when the fishes invoke *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 fish invokes *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 3 times. 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 attempts fail because of an inappropriate approach angle which triggers the *avoiding-fish* response. The male turns away to avoid the collision and tries again.


7 Conclusion

We have demonstrated a framework for behavioral animation featuring an artificial fish model with some astonishing behaviors. These behaviors yield realistic individual and collective motions with minimal intervention from the animator. The easy extensibility of our framework is made most evident by the complex patterns of mating behavior that we have been able to implement to date. Our implementation can run a simulation of 10 fishes, 15 food particles, and 5 static obstacles at about 4 frames/sec (including wireframe


rendering time) on a Silicon Graphics R4400 Indigo2. One of the many exciting directions for future research is to further increase the relevance of our work to the burgeoning field of artificial life [4]. We may be within reach of computational models that imitate the spawning behavior of the female (release of gametes) and the male (fertilization), hence the evolution of new varieties of artificial fishes through simulated sexual reproduction.

References

- [1] H. E. Adler. *Fish Behavior: Why Fishes do What They Do*. T.F.H. Publications, Neptune City, NJ, 1975.
- [2] N. Badler, B. Barsky, and D. Zeltzer, editors. *Making Them Move*. Morgan Kaufmann, San Mateo, CA, 1991.
- [3] V. Braitenberg. *Vehicles, Experiments in Synthetic Psychology*. The MIT Press, Cambridge, MA, 1984.
- [4] S. Levy. *Artificial Life*. Vintage Books, NY, 1992.
- [5] K. Lorenz. *Foundations of Ethology*. Springer-Verlag, New York, 1973.
- [6] P. Maes, editor. *Designing Autonomous Agents*. The MIT Press, Cambridge, MA, 1991.
- [7] D. McFarland. *Animal Behaviour*. Pitman, 1985.
- [8] M. McKenna, S. Pieper, and D. Zeltzer. Control of a virtual actor: The roach. *ACM SIGGRAPH 1990 Symposium on Interactive 3D Graphics*, 24(2):165–174, 1990.
- [9] G. S. P. Miller. The motion dynamics of snakes and worms. *Computer Graphics*, 22(4):169–177, 1988.
- [10] W. Press, B. Flannery, S. Teukolsky, and W. Vetterling. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, Cambridge, England, 1986.
- [11] M. H. Raibert and J. K. Hodgins. Animation of dynamic legged locomotion. *Computer Graphics*, 25(4):349–358, 1991.
- [12] O. Renault, N. Magnenat-Thalmann, and D. Thalmann. A vision-based approach to behavioural animation. *Visualization and Computer Animation*, 1:18–21, 1990.
- [13] C. W. Reynolds. Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21(4):25–34, 1987.
- [14] D. Terzopoulos and K. Waters. Physically-based facial modelling, analysis, and animation. *Visualization and Computer Animation*, 1:73–80, 1990.
- [15] R. E. Thresher. *Reproduction in Reef Fishes*. T.F.H. Publications, Neptune City, NJ, 1984.
- [16] N. Tinbergen. *The Study of Instinct*. Clarendon Press, Oxford, England, 1950.
- [17] X. Tu and D. Terzopoulos. Perceptual modeling for behavioral animation of fishes. In *Proc. 2nd Pacific Conf. on Computer Graphics*, Beijing, China, 1994.
- [18] X. Tu, D. Terzopoulos, and E. Fiume. Go Fish! ACM SIGGRAPH Video Review Issue 91: SIGGRAPH'93 Electronic Theater, 1993.
- [19] P. W. Webb. Form and function in fish swimming. *Scientific American*, 251(1), 1989.
- [20] J. Wilhelms and R. Skinner. A "notion" for interactive behavioral animation control. *IEEE Computer Graphics and Applications*, 10(3):14–22, 1990.
- [21] R. Wilson and J. Q. Wilson. *Watching Fishes*. Harper and Row, New York, 1985.
- [22] D. Zeltzer. Motor control techniques for figure animation. *IEEE Computer Graphics and Application*, 2(9):53–59, 1982.




(a)




(b)

Plate 1: Hook sequence from "Go Fish!"



(a)




(b)

Plate 3: Schooling behaviors.



(a)




(b)

Plate 2: Predator stalking and eating prey.



(a)



(b)

Plate 4: Mating behaviors.