UDC 001.81:519.68:681.32

Autonomy and Artificial Life Forms for Amusement and Use as Agents

●Koichi Murakami

(Manuscript received August 24, 1999)

This paper describes the attributes that can be given to artificial life forms and intelligent agents so that they appear to be alive and posses their own individuality. Also, this paper describes possible architectures for realizing these attributes.

1. Introduction

We started to develop forms of artificial life in 1990, and in 1995 we developed an entertainment software called “TEO – Another Earth” which can communicate with human beings. Then, in 1999 based on this experience, we constructed and released an agent on the Internet.

This paper describes our concept of artificial life, the image of the “living organisms” we intend to create, and some architectures for artificial life.[1] Then, this paper describes various systems we have developed, in particular “TEO – Another Earth,” which we developed in collaboration with Mr. Macoto Tezka, the producer.

2. From agents to artificial life

An agent is a human expert who helps another person do something upon request, and an interface agent is a computer program that operates as an expert.[2] Many researchers are now concentrating their energy on the study of artificial intelligence and agents. For examples of the work being done in this field, see Ref. 2), in which nearly 20 of the foremost researchers in this field (including Minschy, the most prominent figure in the study of artificial intelligence) discuss agents from various points of view.

Researchers originally began studying artificial intelligence for the purpose of understanding human beings and creating human-like robots. The study then branched into several areas, for example, robot engineering, image and voice recognition, and symbol processing for planning and learning. Now, however, these areas are being integrated into a united effort to study human beings.

Agents may possibly change people’s associations with computers. If a person asks an agent on a computer screen to do something, the agent will carry it out to the best of its capability. Agents will come to be regarded as capable secretaries, faithful stewards, obedient servants, and perhaps even close friends.

However, to create these autonomous entities, we will have to better understand the difficult region covered by the theory of communication, where we have not yet found standard solutions even to questions regarding communication between human beings. The questions to be solved include “What makes a person want to communicate with an entity?” and “What makes a person want to continue communicating with an entity?” There are a variety of possible answers – for example, the entity is cute, friendly, interesting, or in the process of growing. We hypothesize that people want to communicate with an entity or agent when they feel that the entity is alive. We cannot prove this hypothesis now; however, based on our implicit knowledge, we can say that human beings have certain special feelings toward living organisms.

An artificial life in a computer cannot be genuinely alive, but it can nevertheless appear to be alive, and it is up to the person who communicates with an artificial life to decide whether it appears alive. We therefore have to solve the fundamental question of what makes people feel as if an artificial life is actually alive. The answer to this question may come from philosophy, biology, or physiology, or perhaps the answer has already been given by Walt Disney and Osamu Tezuka, who is a famous Japanese cartoonist. That is, perhaps the answer is that *animation* gives life to things and characters. After all, we cannot ignore the fact that many of the cartoon characters they created have been "living" in our minds.

3. Architecture of artificial life

The first task of our research into the essential properties of artificial life was to build an autonomous mechanism that behaves adaptively to achieve certain goals on a computer. The autonomous behavior of this mechanism is completely different from the random selections of predetermined behaviors that are performed in the simulated pet systems which are enjoying recent popularity. The autonomous behavior of an artificial life is dependent on factors in the external environment (e.g., weather, time of day) of the user and other living organisms and the artificial life's internal conditions (e.g., degree of hunger, emotions, and degree of familiarity with the user). If we use a conventional programming paradigm to realize autonomous behavior, we would have to describe all the conditions of an enormous external environment in a list. This, however, is not a practical approach. In fact, robot researchers have faced the same problem.

To achieve its goal, a robot must determine its actions by itself in a dynamic environment that includes moving objects such as people and other robots without hitting them. Therefore, the architectures developed by robot researchers are planning systems which react to conditions. **Figure 1** shows one such architecture we have developed at Fujitsu. These architectures consist

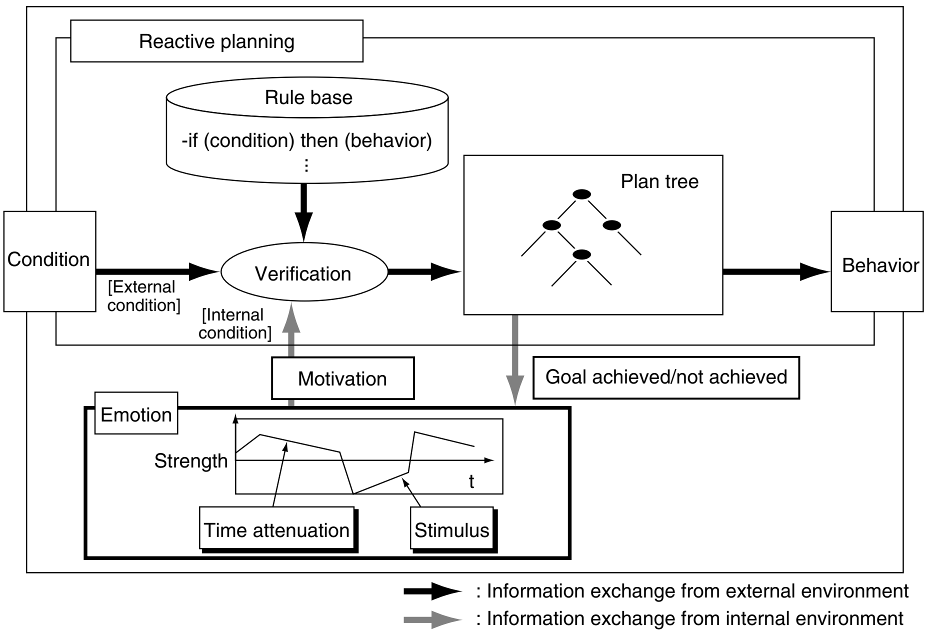

Figure 1
Autonomous architecture.

of three modules: one for recognizing conditions in the external world, one for judgment (thinking) based on this recognition, and one for executing the actions to be taken. The main feature of these architectures is that the three modules are linked with each other in a loop in realtime, so that when a robot takes an action in the external world, it quickly recognizes the effects of the action and then makes the appropriate judgement or judgments. Because of this loop the robot maintains an internal analog of the external world.

In the judgment module, which is the core of these architectures, behavior is described as a set of rules for actions. The relationships between actions in this module and how they influence the robot's behavior are not predictable but depend on conditions in the external world. In other words, individual actions are only meaningful in terms of external conditions.

A robot having such an architecture, however, is not a living organism because it is always given a goal in advance. Let us take the example of a cleaning robot that moves around and collects dust and trash while making way for human beings, cats, dogs and other obstacles. Such a cleaning robot can be made with an extremely high level of technology and could make enormous contributions to society. However, it will only act for the purpose of cleaning and will still be just a machine.

A living organism can have simultaneous, important goals, for example, self-preservation and reproduction. A living organism might become confused when some of these goals conflict with each other and find which is the most important goal for that situation. It eats when it is hungry, looks for a friend when it feels lonely, and sleeps when it is tired. When people see or recognize such behavior in a living organism, they try to deduce the living organism's state of mind from its behavior. The essential points here are that the goals of a living organism are not dictated by an external source but are produced autonomously according to motivations internal to the living organism and that when a living organism demonstrates this autonomy we are likely to regard the organism as possessing a degree of consciousness.

However, an autonomous mechanism cannot be created simply by designing an artificial life that reacts to the external environment; it can only be created when an artificial life is given much more profound internal conditions. To achieve this, we developed an emotion module with parameters for physiological states such as hunger and fatigue and feelings such as sorrow and joy (see Figure 1). These parameters quantitatively change depending on the artificial life's own actions, its recognition of external conditions, and the degree to which it achieves its goals. As a result, a change in circumstances can produce a new goal inside the artificial life.

Consider, for example, the feelings of like and dislike and assume that an artificial life "likes" an object. If the object is in front of the artificial life, the simple goal of "getting the object" is produced. If the object is at a distance, the goals of "finding," "reaching," and "getting" the object are produced. The artificial life then takes the necessary actions to attain the produced goal or goals. While executing these actions, the artificial life monitors and reacts to changes in the environment because the recognition, judgment, and action modules continue operating in the loop described above. If someone brings the target object close to the artificial life, it will have achieved its goal without the need to move to the object's original location, so it stops moving. Then, because its goal has been achieved, the artificial life can autonomously take a new action. For example, the artificial life may be pleased and begin laughing or dancing or begin to "love" the person who brought the object and follow that person around. Thus, we built an artificial life mechanism which behaves according to its own goals and appears as if it has a mind.

Our first artificial lives made their debut in

Figure 2
"TEO – Another Earth."

CG Sapporo in 1990. The presentation contained a shark and three jellyfish. The shark chased the jellyfish when it was hungry, and the jellyfish gathered together to play with human visitors to their world. Wearing data gloves, people could have a type of conversation with the artificial organisms through gestures. In 1992, based on our first attempts, we developed another world of artificial life named "Charlotte," in which 10 artificial organisms lived. These artificial organisms had the autonomy mechanism described above, which made them behave with an individuality. Some of them had friendly relations with each other while some of them struggled with each other, and they formed groups similar to the groups that are formed in a society of monkeys.

4. TEO – Another Earth

Unfortunately, we lost interest in communicating with these artificial organisms after about 20 minutes. Then, we made an improved version called "TEO – Another Earth" (see **Figure 2**). TEO is an imaginary planet on which various imaginary organisms that we can communicate with are living.

The first artificial life we developed for TEO is an imaginary animal called "FinFin" that looks like a cross between a dolphin and a bird. Mr. Macoto Tezka, an artist with a profound knowledge of living organisms, helped us to develop TEO in the capacity of producer. Through our discussions with him, we determined that we would not create a character or artificial organism that simply flatters and plays with human beings. Instead, we decided to create an artificial organism that has its own emotions and autonomously lives its own life.

We succeeded in realizing some degree of autonomy in TEO. The autonomous architecture which decides FinFin's behavior also controls the planning of its strategic behavior. If FinFin is hungry and finds food nearby, FinFin will eat it. If FinFin is playing with somebody but does not find any food, it will ask the person for food. If there is nobody, FinFin will go to another place to eat something. FinFin is an artificial life that in certain ways appears as intelligent as a dog or cat.

A living organism reacts to the environment, and if its enemy appears, it will run away. FinFin also has a mechanism to take action in an emergency. A typical example of an emergency for

FinFin is a person speaking loudly. If a person playing with FinFin speaks loudly, FinFin detects the loud voice through a sensor attached to the computer and immediately runs away. FinFin will then increase its wariness and not show up for a while.

FinFin has the emotions described before and behaves according to its emotions as if it had its own will. Such behavior sometimes manifests as disobedience to the person playing with FinFin. However, the person can guess what is on FinFin's mind from the context of its behavior and then act appropriately. This can be the first step of communication. FinFin sometimes listens to people but does not fully obey instructions. This function is the fruit of FinFin's autonomous architecture.

As explained above, we developed functions for reacting to circumstances and motivating behavior using models of emotions and a planning system based on a set of rules. We can predict most but not all of FinFin's behavior. We believe that FinFin has the changeability, complexity, and dependency of a living organism and is advanced enough for us to have empathy with it.

Of the people who have bought TEO, 60 percent are female, and many of the buyers also bought a personal computer at the same time so they could run TEO. We also found that many TEO users have a strong empathy with FinFin and enjoy communicating with it. TEO users can be classified into the following four groups according to the way in which they enjoy TEO.

1) Users who mainly watch the dynamic changes in the TEO environment

There are many users, chiefly middle-aged men, who mostly enjoy watching the changes in the imaginary world. On TEO, a day consists of 24 hours and a year consists of 365 days. The natural environment and lives of living organisms alter as time goes by. When night falls, stars and a moon appear and the living organisms sleep. In the morning, the living organisms wake up, eat, and sing lively. Watching these changes in the environment, these users might be wishing they could enter the jungle-like world of TEO themselves.

2) Users who enjoy TEO as a game

In an exhibition, we saw five school children rush to the demonstration TEO machine and begin using it exclusively. While we were observing them, they competed with each other for FinFin's attention. Girls fed FinFin and called its name tenderly; boys, on the other hand, threatened FinFin loudly. Playing these types of games is a popular way of enjoying TEO among school children.

3) Users who simply enjoy communicating with FinFin

In the exhibition, we observed that preschool children enjoyed simply communicating with FinFin. What was remarkable was their conversations with their mothers while they were playing with FinFin. The children and their mothers actively talked to each other about FinFin's every action. This suggest that an artificial life can serve as a medium of communication between human beings.

4) Users who communicate affectionately with FinFin

Some users, mainly unmarried young people, have sent messages to our TEO home page asking such questions as "How do you turn off the computer when FinFin is eating in front of you?," "I am anxious about what FinFin is doing while I am absent," and "How can I feed FinFin while I am away?" They seem to have been taking care of and communicating with FinFin with affection for over a year.

From our observations of TEO users, we conclude that an artificial life communicating closely with human beings and behaving autonomously can attract people in various ways. Since communication requires an individuality, an individuality is a requirement of any artificial life that functions as a human being's communication partner. We think that TEO has proved this. With the users of the fourth group described above, daily pet-like communication is particularly important. It is therefore important that users are able to make this type of daily communication.

5. Development into an agent

The agent we want to construct resembles Tin Man or Scarecrow from the movie "The Wizard of Oz." The general image of an agent may be something like a secretary. Our image of an agent, based on the experience we gained in developing TEO, is described below.

5.1 User's feelings toward an agent

To encourage a user to communicate with an agent for a long time, functional perfection alone is not enough. We have to design an agent that is in some way addictive. TEO users initially felt that FinFin was pretty, but then became more attracted by FinFin's complex, autonomous behavior. We therefore want to add additional features to an agent so that users continue to find it interesting.

5.2 Duration of user's communication with an agent

If an agent is functionally perfect but otherwise uninteresting, users will communicate with it only for a short time. For example, if a user asks such an agent to book a train ticket for Tokyo, the agent might be able to complete the task within a minute. Also, in the real world, for example, managers do not talk all day with their secretaries, even when they want their secretaries to perform a large number of tasks.

5.3 Functional requirements

There remains the following problem regarding the functional requirements of an agent. Usually, a secretary does not have expert knowledge of a narrow domain, which is the strong point of artificial intelligence, but a general knowledge of a wide domain. However, in general, it is difficult to create an artificial intelligence that has a general knowledge of a wide domain; therefore, an agent requires an interactive function so that the users can help the agent perform a task by conversing with it. The problem here is that currently it is technologically difficult to equip an agent with an interactive function. So far, our agent has the following functions:

1) Guidance to home pages

By selecting an automatic function, the agent can guide the user to various points on a home page, so the user can look at or read through a home page as if he or she were on a bus tour. This is convenient because it allows the user to overview the contents of a home page without needing to make extra operations.

2) Management of personal information

For example, if our agent was used as the agent for a home page about pets, it could occasionally notify the user about information such as the pet's name, date of birth, and when to buy food for the pet. If the pet's birthday is approaching, the agent could inform the user. If the user's stock of pet food is becoming short, the agent could advise the user to buy some more (through mail order).

3) Playing with the agent

As described above, we think this is an important function. A function for feeding a user's artificial pet, for example, is significant because the user can feel like the pet's owner or one of its fiends. If the artificial pet eats its food every time the user gives it and takes no action even when it has not been fed, the artificial pet will not seem real. Therefore, an artificial pet must ask the user for food when it feels hungry. Based on our experience with TEO, we included this function in our agent. We also gave the agent functions for play that are specific to the World Wide Web because agents use the web to communicate with users. To allow the user to be actively engaged while accessing home pages, there is a game-like function to let the user and the agent jointly look for the pet's favorite food (e.g., bones) hidden in home pages. Finding and accumulating this food will give the user a sense of achievement.

4) Acquisition of users' profiles

Regarding the application of our agent to electronic commerce and other on-line activities, a future issue will be how to acquire the dynamic profiles of users. In various experiments, we have already attempted to acquire the implicit, dynamic profiles of users through the interactive communications made between the agent and users.

5. Conclusion

This paper looked at various topics regarding artificial life forms for human amusement and assistance, for example, the design of their architectures and the importance of the appearance of autonomy. Then, this paper described an imaginary world we have created called "TEO" and an artificial life form called "FinFin" which inhabits it. Lastly, this paper looked at various aspects regarding the design of intelligent interface agents for use on the World Wide Web.

We will widely incorporate the experience we have gained from creating these imaginary worlds and artificial life forms into agents and into human interfaces. We will also release an authoring system for Internet service providers who intend to create new products based on these technologies.

Acknowledgments

The ideas put forward in this paper are the valuable products of research done at the Computerized Agent Project of Fujitsu Limited. I gratefully thank the project staff for their challenges to do what has not yet been attempted. I also thank Mr. Macoto Tezka and the other artists who helped us create TEO.

References

1) K. Murakami: Interactive TV, Board of Industrial Investigation. Conversation between Human Beings and Computers (joint work).
2) Intelligent Agent. *Communication of ACM*, **37**, 7, 1994.

Koichi Murakami received the M.Sc. degree in Nuclear Engineering from Hokkaido University, Sapporo, Japan in 1980. He joined Fujitsu Laboratories Ltd., Kawasaki, Japan in 1980 and was engaged in research of computer graphics, artificial intelligence, and parallel computers. In 1991 he was a visiting scholar of Carnegie Melon University, and in 1992 he was awarded the Sakai Special Award. Currently, he is engaged in research and development of artificial life and human interfaces. He is a member of the Information Processing Society of Japan (IPSJ).