FEATURE 16 September 1995

Oh look he's brought me a present

SILAS is a friendly-looking golden dog with floppy ears and a pointy tail. If you hold out your hand, Silas will come up to you wagging his tail and sit down, just like a real dog. He will even bring you a ball and try to get you to play.

But Silas is not real. He is made from computer-generated shapes and looks a little as though he escaped from a Disney cartoon. Silas is one of a small but growing number of computer programs that can mimic the behaviour of friends, pets and other objects of human affection. Among them are Julia, with whom you can have long, if strange, conversations, Phink, an unpredictable dolphin-like creature, and Neuro-Baby, a digital infant who can analyse and react appropriately to your moods.

Building "relationships" with these digital creatures is, of course, rather an unusual process. With Silas, for example, you have to stand in front of a video camera looking at a projection screen on which you see a picture of yourself, your room and Silas, added in by the computer program. Of course, these are all virtual friends; robot friends with "real" bodies are still a long way off. The first clumsy autonomous robots – among them Skimer – are only now emerging from the laboratory. One day, they may turn into the mechanical servants beloved of early science fiction writers.

For the more immediate future, the scientists and artists who are creating digital companions want to explore what is needed to make these computer-generated objects more friendly and engaging to humans. That research is proving that we are often surprisingly willing to read intelligence and intention into the creations of computer programs.

Attempts to build such software started back in the early days of artificial intelligence research. The inspiration to write programs that could provide sufficiently human-like communication to fool other humans came partly from the famous "Imitation Game" set by British mathematician Alan Turing. His view was that if a computer program mimicked a human so effectively that a real human could not tell whether they were interacting with another human or a computer, then the program was effectively a human being.

Early approaches could tackle only text-based interactions because computer power was far too small to generate a real-time visual image as complex as a dog. Eliza, developed in 1965 by Joseph Weizenbaum at the Massachusetts Institute of Technology, was one of the most famous of the conversational programs, imitating the questioning of a therapist. The program scanned the user's punched input for certain words and phrases, then plugged these phrases into sentences from its database. The resulting simple, open-ended questions seemed to come from a human who was listening to what the user was saying:

Human: I feel unhappy.

Eliza: Is it because you feel unhappy that you came to see me?

This was a clever strategy: by adopting a style of discourse designed to conceal the speaker's personality, it could also conceal the fact that the speaker had no personality to conceal.

Parry, developed a decade later, approached the problem from the opposite perspective. The program imitated a paranoid, and the user took the role of the therapist. A typical exchange looked like this:

Human: How did you come to be in the hospital?

Parry: I am upset.

Human: How long have you been in the hospital?

Parry: About a week.

Human: In what way are you upset?

Parry: People make me nervous.

Human: In what way do they make you nervous?

Parry: People give me a funny look sometimes.

Parry's imitation of a paranoid encouraged the user to accept its non sequiturs and disconnected responses; if the program was unable to find anything in the user's input that it could understand, it could simply blurt out something at random without seeming to be acting out of character.

Research on text-based question and answer programs has developed rapidly since then. Julia, developed by Michael Mauldin at Carnegie Mellon University's Center for Machine Translation in Pittsburgh, Pennsylvania, has come much closer to imitating human conversation. Julia imitates a normal person – one that, by coincidence, seems to want to talk about pets:

Human: Do you have a pet at home?

Julia: I'd rather have a real pet, like a cat. Cats are my favorite, because you can pet cats, and they're easy to keep.

Human: Do you know anything about miniature schnauzers?

Julia: I don't know anything about miniature schnauzers.

Human: Well, I have one. We call him Pucci. He is a crazy cat – have you ever had a cat?

Julia: What a coincidence! I have two cats.

Julia uses many of Eliza's and Parry's tricks. For instance, since asking quetions is much easier than answering them, Julia asks a lot of questions, like Eliza. And like Parry, Julia frequently takes control of the conversation by changing the subject. But Julia adds a number of new tricks. It appears to have a more vivid personality than its predecessors, because its database includes statements that are controversial or funny, such as "Guinea pigs are about as smart as cucumbers". The database also includes fragments of more extended conversational threads that can make the program seem coherent if the user is fooled into playing along with them.

Julia cannot think. But as long as you don't try to move beyond pleasant small talk about pets, you stand a good chance of being fooled.

The key to success is to isolate the characteristics that make interactions seem real. We are often willing to explain away meaningless behaviour as idiosyncratic, and to ascribe intention where none exists, provided we have reason to believe we are dealing with an animate object.

It's a dog's life

Silas the dog builds on these insights. "The research goal behind Silas is to understand how you can build an autonomous creature, like a dog, that seems to do the right thing over time," says Bruce Blumberg, Silas' creator and a PhD student at the Massachusetts Institute of Technology's Media Lab. People readily respond to Silas as though he were real, creating explanations for his behaviour the way they would with a real dog. "From the user's perspective," says Blumberg, "there's a sentient, intentional being there."

The willingness of humans to see complex emotions in animals is well known to ethologists. "People often ascribe feelings to a real dog that are over and above what they really would admit, if pushed, they believe the dog really feels," says June McNicholas, a research fellow at the University of Warwick who has worked extensively on the bond between humans and animals. "If you've had a bad day at work," she says, "your dog may seem to respond to this. And you'll say to yourself 'he knows I had a bad day at work'. But you know that the dog doesn't know you had a bad day at work. All he knows is that you are moving and acting differently than you usually do. And you know this."

Reacting to Silas involves similar rationalisations. For instance, Silas is interested in moving objects that are close to him. If you reach out to pat him, your hand will be both moving and close, and Silas will watch it. Silas doesn't have any way of knowing whether you're touching him or not, but because he's watching your hand so intently he "appears" as though he is responding to being patted.

Users tend to explain his response in those terms ("I'm patting him, he likes it") and the explanation has predictive value (whenever the user tries to pat the dog, the response is the same) and so, in the user's mind, Silas likes to be patted.

People find Silas fascinating partly because they enjoy trying to explain what he is doing. "If a creature behaves exactly the same every time," says Blumberg, "that's not very interesting. It turns into a robot. On the other hand, if it's totally unpredictable, then it seems random, and it's hard for the user to develop an explanation with predictive value. The optimal place is where there's just enough surprise that you're constantly coming up with new explanations that make sense. I think that this is why we like having ‘real’ pets.”

Blumberg hopes that his work will soon be built into video games. By making creatures in a game more responsive and by adding autonomous and unpredictable behaviours, they may seem more alive and more interesting. Indeed, Japanese game designers are already experimenting with exactly these components.

Flying dolphin

The Believable Agent Project is Fujitsu’s foray into the world of virtual creatures. Software designer Makoto Tezuka has created a CD-ROM world called “The Other Earth” (TEO), where Phink, a flying dolphin-like creature, lives.

To help you communicate with Phink, Fujitsu provides a “TEO antenna”, a combined infra-red sensor and microphone that connects to your computer and can pick up voice and movement, and a whistle which can be blown to attract Phink if it isn’t on screen. But Phink will not necessarily respond.

The idea is to give Phink the quality of a wild creature, which can only be “tamed” if you are really nice to it, and don’t try and keep Phink as a pet. Users can’t tell in advance what mood Phink will be in. They can’t even tell what image they will see when returning to TEO after a couple of days. Time passes on TEO, and the world changes, even when the software is inactive.

Whether TEO, which has not yet been released for sale, is a success remains to be seen. At best, however, TEO’s ambitions are limited to providing a more engaging game environment. Elsewhere in Japan, video artist Naoko Tosa has much bigger goals. She wants her Neuro-Baby to “live and communicate with modern urban people like ourselves, people who are overwhelmed, if not tortured by the relentless flow of information, and whose peace of mind can only be found in momentary human pleasures”.

Tosa is visiting artist at the ATR Media Integration Communications Research Laboratories near Kyoto and is developing her latest Neuro-Baby with researchers from the University of Tokyo. The baby’s present home is a run-down laboratory, packed with computers, young programmers and cabling in a basement of the decaying Institute of Industrial Sciences in Tokyo.

Where other researchers have tried to add interactive and unpredictable behaviour, Tosa has gone for emotion. On screen, Neuro-Baby appears deceptively simple – just an animated cartoon infant face with big round eyes, expressive eyebrows, and a small set of utterances provided by a voice synthesiser. But Neuro-Baby is more than it seems: it can analyse voice stress patterns to determine your mood and to respond accordingly. The voice stress analysis is provided by a trainable neural network computer.

Talk to the baby and it responds as if it is sharing the same emotion. The baby’s reaction is always interesting, often disconcerting. Its emotions, says Tosa, run from “cute” to “stubborn and introverted” and on through “conceited and perverse” to the “crazy type with violent emotional ups and downs”. If you’re feeling relaxed and peaceful, Neuro-Baby can be sugarysweet, asking you with a cute smile “Do you live around here?” Ignore Neuro-Baby, and it can start to irritate as it tries to get your attention, asking “Do you know any tongue-twisters?” It knows lots and it will tell you them all. Snap at Neuro-Baby and its face grows angry and red, and its head jerks around the screen as it shouts “Stop it! Stop it!” at you. To make interaction more real, Neuro-Baby maintains eye contact. Two video cameras track the eyeballs of both the speaker and the baby so they are always locked onto each other.

An earlier version of Neuro-Baby was exhibited at the SIGGRAPH exhibition in 1993. There, Tosa found that when Americans interacted with Neuro-Baby, it started to act as though it were emotionally unstable. Americans were just too much for it. “It was educated with Japanese monotonous voices,” says Tosa. “American intonation was too strong for it.” Now she has made separate emotional models for Japanese and American users, and at this year’s SIGGRAPH she tried using Neuro-Baby as a cross-cultural emotional communication tool.

When a Japanese and an American communicate through Neuro-Baby, the Japanese sees the American’s emotions toned down to Japanese levels; the American sees the Japanese participant’s emotions expressed by Neuro-Baby in the more extreme American way. Neuro-Baby could eventually act as an electronic agent which translates emotions as well as language.

This is ambitious, but Tosa has already achieved considerable fame in Japan for her work in experimental film and video which, she explains, are on the boundary between “the visible and the invisible”. She see enormous potential for individuals to create digital companions. “I wanted to create something like my alter ego,” she says. “Some organism which is also like my closest friend or relative.” So Neuro-Baby is also a means of personal expression. “People express their dreams in the media at hand, such as novels, films and drawings,” she explains. Neuro-Baby uses contemporary media to reach the same ends.

Tiny steps

Behind Neuro-Baby’s artless reactions there is much thought and a lot of computer power. She is now adding a hand which you can shake to the latest model. Sensors monitor the form of your handshake and feed additional information into the computer’s model of your emotional state. This is Neuro-Baby’s first tiny step out of the computer. But there is a very long way to go before it, or Silas, can become “real”.

Bringing artificial pets into the outside world is far more complex than creating images on a screen. This is why robotic pets are many steps behind virtual ones. Will Wright, the inventor of SimCity, the interactive computer game that allows users to create their own urban dreams, has a passion for robots. But, as he explains, out in the real world, a robot pet can’t see an object unless it actually detects and interprets the light reflecting off it. “Edge parsing is fairly computationally intensive,” he says, “and it’s very hard to do it in real time.”

Indeed, Silas’s ability to identify where the user is standing depends on the much easier technique of “background subtraction” – the computer knows what the room looks like when it’s empty. When a person walks into the scene, Silas’s image-processing software compares the stored image of the empty room with the current scene – what remains must be the person.

This technique depends on an unchanging background, though. Start moving furniture around, or bring more people into the room, and Silas is likely to get confused. And while Silas is able to process and parse the video image in real time, these tasks are handled by a Silicon Graphics Iris workstation, which is rather more computing power than will fit in a mobile, pet-sized box.

Skimer the robot illustrates these constraints. Kino Coursey, of Daxtron Laboratories in Fort Worth, and the developer of Skimer's software, says that the robot can visually identify objects and be trained to follow them around. To do this, one trainer drags a chair, say, while the other uses a joystick to instruct Skimer to move in pursuit of the chair.

Skimer builds a network of associations between the images that it is seeing and the commands it is receiving. For example, when the chair moves out of the robot's field of view, moving from right to left, the trainer commands Skimer to turn to the left. Skimer remembers the sequence of images and the commands associated with them. Once Skimer has been trained, it will follow the left-turn command whenever it sees a chair moving across its vision to the left. The result, says Coursey, is that "you can drag a chair around in front of it and he'll follow it anywhere".

But while Skimer does this very well, it can't do anything else. And it's a pretty hefty contraption, cobbled together from a child's six-wheeled riding toy, a camcorder, a computer, and 18 kilos of batteries. Skimer, as Coursey puts it, "definitely belongs in the back yard. You wouldn't want him in the house".

Pet care

If Skimer could be made small and agile enough, could it replace the family dog or cat? Erika Friedmann, a specialist in pets and health in Brooklyn College's Health and Nutritional Science Department, acknowledges that the autonomy of artificial pets might attract people the way real pets do. "People like pets because they don't have to make an effort to get positive feedback from another being," she says. "Your pet makes some kind of acknowledgment that you're there and entices you to interact with it."

But pets provide people with far more than unstructured entertainment, says Friedmann. They give us someone to care for, someone to touch and fondle. They provide a reason for exercise, a feeling of safety.

McNicholas thinks that this is the fundamental weakness of the artificial dog. "What kind of level of care could someone give a virtual pet?" she asks. "A lot of the closeness in a relationship with a pet is based on how dependent the pet is on you."

Even the most complex artificial pet only engages people on one of the many levels that real pets do. Though you can turn off Silas or Neuro-Baby, these aren't "deaths". You can always turn them on again. Artificial pets don't depend on your care to keep them alive and well. And, as McNicholas says, "if an artificial pet doesn't depend on you, you don't feel needed". Perhaps that justifies Tosa's view of the potential of digital companions: they may not be able to replace pets, but they could give people a new way of expressing themselves.

By **Robert Rossney**

Freelance Journalist, Berkeley, California;
Reporting From Japan was provided by Shinichi Endo, Technology Editor, Newton, Japanese Monthly Science Magazine

Magazine issue 1995 published 16 September 1995