

Social Robots

Scientists are creating machines that can express emotion, talk, and even learn.

Will robots become man's new best friend?


by Hemma Mistry

Imagine this scenario: a man and a woman converse with a third person, a judge, through messages on a computer. The judge does not know which person is the man and which is the woman, and the man and woman try to mislead the judge about his or her identity. The judge must determine which is the man and which is the woman using only their conversations through the computer. Now imagine either the man or woman is replaced by an artificially intelligent machine. Again, the judge must decide which is which only through the conversation, but if the machine is able to fool the judge into believing it is human, can it be called intelligent? This situation is known as the Turing test, proposed by mathematician and computer science pioneer Alan Turing in 1950 (Flood 1996).

If the machine can pass this test, it could demonstrate that the machine possesses a simple theory of mind, or the ability to know that other people have different thoughts and beliefs than one's own. The question of whether machines will ever be able to think like humans has fueled debate amongst scientists and philosophers for decades, and the question becomes more important as machines become increasingly able to mimic human behavior.

Robots that Learn

Although the possibility of artificial consciousness is still debated, computer scientists and robotocists have been developing artificial beings that at least mimic the characteristics of human consciousness. In the last fifteen years, huge progress has been made in the development of robots in terms of social interaction. Robots such as these are known as "Social Robots" because of their ability to interact with people in a humanlike manner. Social robots are autonomous, meaning that they can make some decisions independently of their human controllers. They can also be characterized as situated (they can sense and interact with their environ-




Babybot learns through trial-and-error to grasp objects with its arm.

ment) and embodied (they have physical bodies that interact with their environment) (Henig 2007).

Before research into social robots became popular, it was thought that the highest kind of robot intelligence was abstract intelligence, which involves, for example, solving complicated mathematical problems. Now design goals focus more on creating robots that learn like humans: through trial and error. Instead of being pre-programmed to take an action, social robots of today learn to perform an action through experiment and mimicry, just as children learn. This action-based approach to artificial intelligence is known as embodied intelligence (Henig 2007).

An early example of a robot designed with embodied intelli-

Instead of being pre-programmed to take an action, social robots of today learn to perform an action through experiment and mimicry, just as children learn.



<http://www.ai.mit.edu/projects/ntt/projects/NTT9904-01/kismet-images/kismet-interesting.jpg>

Kismet shows expressions of interest, boredom, and fear due to social interaction.

gence is Cog, which was built at M.I.T.'s Humanoid Robotics lab from 1993 to 2004. Among other things, Cog could learn to adjust the position of its arms according to the weight of objects placed in its hands (Henig 2007).

More recently built robots are engineered with a slightly more sophisticated ability to learn. Designed to learn like a child, Babybot, the creation of researchers at the Laboratory of Integrated Advanced Robotics (LIRA-Lab), learns to interact with objects through experimentation. The robot features a torso with one arm and a hand with fingers, with which it learns to grasp objects through trial and error. The Babybot project, started in 1996, has led to the ongoing 2004 Robotcub project, an attempt to build a humanoid robot about the size of a two year old child. Robotcub will have the abilities of Babybot but also the ability to learn how to crawl and maybe even to walk (Simonite 2006).

The robots' ability to learn like a human child is an exciting development in robotics, but how much can robots actually learn on their own? Rodney Brooks, designer of Cog and former director of M.I.T.'s Computer Science and Artificial Intelligence Laboratory, explains, "They can only learn certain things, just like a rat can learn only certain things and a chimpanzee can only learn certain things and even [you] can only learn certain things" (Henig 2007). Ultimately, the robots must first be programmed to learn an action by its human designers, which is a huge limitation on the potential of artificial intelligence.

Socially Intelligent Robots

Along with the ability to learn like humans, social robots can mimic the behaviors of human social interaction. By following certain social cues, these robots encourage humans to interact with them comfortably and can even cause people to perceive them as living creatures. Kismet, a robot built in 1996 by Cynthia Breazeal at M.I.T., was made with human interaction in mind. Kismet comprised of a head with eyes, eyebrows, lips, and ears which all moved to create facial expressions (Henig 2007). It could

express emotions depending on whether or not its pre-programmed desire for stimulation was being met. For example, human interaction caused the robot to assume an expression of happiness, over-stimulation caused its eyebrows to rise in an expression of fear, and under-stimulation caused an expression of boredom with droopy eyes. Kismet was also capable of holding what Breazeal calls "proto-conversations." Although it could only make babbling noises, the robot was able to hold vocal exchanges with people by using social cues. These included leaning forward and raising its eyebrows to prompt the other person to speak and leaning back when it intended to speak (M.I.T. Artificial Intelligence Lab).

Breazeal is currently building Leonardo, or Leo, a robot whose exterior was designed at the Stan Winston Studio by Hollywood animatronics experts, giving it an unusually life-like appearance. Leo has been programmed to pass the false belief test, a simple test which, like the



<http://blogs.spectrum.ieee.org/automaton/Leonardo%20robot.JPG>

The unusually life-like Leo was designed by Hollywood animatronics experts at the Stan Winston Studio.

Turing test, demonstrates theory of mind in the subject. This test can usually be passed by children at the age of four or five, and it shows a child's knowledge that other people have different thoughts than his or her own (Henig 2007). The test proceeds as follows: Leo observes two people in the room put an object into box 1. After one person leaves the room, the second person switches the object into box 2. The second person then leaves the room and the first person returns. When the first person attempts to open box 1, Leo understands he wants the object and opens a separate box containing the same object. Using sensors, the robot is able to keep track of each person's beliefs about the location of the object, and by comparing it with its own knowledge, can understand what the first person wants (Henig 2007).

The Debate over Artificial Intelligence

While robots such as Kismet can make life-like expressions of emotion, can the machine itself ever really experience an emotion? Kismet's human-like emotional expressions caused the students and researchers who interacted with it to unconsciously treat it as living, even though they knew it was a robot. So are these perceived emotions in any way "real"? Breazeal suggests that while robots can never truly experience human emotions, the question instead should be, "What are the emotions that are genuine for the robot?" (Henig 2007). Just as animals can experience emotion within their conscious realms, robots could have unique robot emotions that they experience in their robot consciousness. Philosopher Daniel Dennett proposes a similar idea in his paper, *Consciousness in Human and Robot Minds*. Dennett refutes the claim that "nothing could properly 'matter' to an artificial intelligence, and mattering...is crucial to consciousness" by comparing the pleasures and pains of simpler organisms to those of robots (Dennett 1994).

Examining the way the human brain functions also leads to a further argument for robot consciousness. In a sense, the brain acts like a complex machine composed of neurons that work and interact according to physical laws. Consider a thought experiment proposed by computer scientist Jaron Lanier: one neuron in the human brain is replaced by an artificial neuron with the ability to perform all the same functions as the original. The brain will function exactly as it did before. Now, the remaining neurons in the brain

are successively replaced by artificial neurons. It would appear that the original conscious state would be maintained (Lanier 2006). Similarly, in a New York Times article, Rodney Brooks argues, "A robot's level of a feeling like sadness could be set as a number in computer code... But isn't a human's level of sadness basically a number, too, just a number of the amounts of various neurochemicals circulating in the brain? Why should a robot's numbers be any less authentic than a human's?" (Henig 2007). If robots are programmed to have certain desires and goals, how do these differ from the desires and goals felt by people?

Human-Robot Interaction

Critics such as John Searle, professor of philosophy at UC Berkeley, still contend that any intelligence a robot may appear to possess is still just an imitation. In his 1980 paper, *Minds, Brains, and Programs*, Searle argues that a machine has no real understanding of what it is programmed to do. He states, "Such intentionality as computers appear to have is solely in the minds of those who program them and those who use them, those who send in the input and those who interpret the output" (Flood 1996). This argument raises an interesting point: a robot's display of consciousness may be completely artificial, but it can reveal a wealth of information about human consciousness. Today's robots are far from mimicking the complexity of the human mind, but even in their current primitive state, social robots elicit surprising responses in the humans who interact with them. Studying human-robot interaction can provide tremendous insight into human consciousness and human social interaction.

Developers of social robots are finding interesting and unexpected responses to their life-like machines. For example, Kismet's facial expressions and ability to converse caused emotional responses in people who interacted with it, even when they knew that it was simply a machine. Because the robot was able to mimic what humans interpret as happiness, people who interacted with it derived pleasure out of making Kismet "happy" (Henig 2007). The study of human-robot interaction is growing into an exciting new discipline that blends psychology, cognitive science, and robotics. Researcher Frank Heger at Bielefeld University is studying human-robot interaction by scanning the brain of test sub-

"A robot's level of a feeling like sadness could be set as a number in computer code... But isn't a human's level of sadness basically a number, too, just a number of the amounts of various neurochemicals circulating in the brain? Why should a robot's numbers be any less authentic than a human's?" --Rodney Brooks

jects as they interact with robots. Heger found that the more human-like the robot, the more the human subjects responded to it as if it were a conscious being (Marks 2008). Other studies, such as one conducted by Paul Schermerhorn at Indiana University, are showing gender differences in robot interaction. His study demonstrated that men were more inclined to engage socially with a robot, whereas women felt emotionally detached from the robots and described them as machine-like (Marks 2008).

The study of robot-human interaction may lead to the integration of sociable robots into our daily lives. Current research examines how machines can use social interaction to improve their function. For instance, M.I.T. graduate student Cory Kidd is studying the effectiveness of a robot designed to be an exercise coach. The robot gives reinforcement and encouragement based on the subject's daily eating and exercise habits (Henig 2007). Studies such as Kidd's are determining what kinds of social interaction elicit positive responses from people. It is easy to imagine robots such as this one becoming integrated into our homes to improve our lives.

While social robots may become an important part of human life, it seems likely that they will remain more like friendly machines than conscious beings, at least for the time being. The Turing test, proposed up by Alan Turing over fifty years ago, has yet to be passed by a machine. In fact, a competition held by Howard Loebner in 1991 offered a prize of \$100,000 for a machine that could trick ten judges in three hours into thinking it was human. The competition was held for a few years until being cancelled; there was no machine able to win the prize (Flood 1996). Artificial intelligence remains a distant possibility, but the new study of robot-human interaction holds much promise for incorporating robots into our lives.

REFERENCES

- Dennet, Daniel C. 1994. Consciousness in Human and Robot Minds. Paper presented at IAS Symposium on Cognition, Computation and Consciousness, Kyoto, September 1-3, 1994.
- Flood, Gary. 1996. If only they could think. *New Scientist* 13 January 1996.
- Henig, Robin Martinez. 2007. The Real Transformers. *The New York Times* July 29, 2007. <http://www.nytimes.com> (accessed February 28, 2008).
- Lanier, Jaron. 2006. Jaron's World: The Soul of the Machine. *Discover*. 25 June 2006.
- Marks, Paul. 2008. The rise of the emotional robot. *New Scientist* 5 April 2008.
- M.I.T. Artificial Intelligence Lab. Humanoid Robotics Group. <http://www.ai.mit.edu/projects/humanoid-robotics-group> (accessed April 5, 2008).
- Simonite, Tom. 2006. 'Baby' robot learns like a human. *New Scientist* 5 May 2006.