

The emotional robot

Cognitive computing and the quest for artificial intelligence

Humans and computers are fundamentally different. Apart from the ungainly exterior of the latter, the main difference is the way in which they process information. Computers follow the instructions of programmed algorithms and user input, whereas the human brain processes information in a nonlinear way with often-unexpected results—which explains much of human inventiveness and creativity. Yet, it might not be too long until computers ‘evolve’ to emulate human cognition—with some help from their human masters, of course.

...a dream that is as old as the golem of Jewish folklore [...] intelligent robots that understand human emotions...

Researchers from fields as diverse as computer science, mathematics, neuroscience, kinematics and cognitive science are getting closer to creating computers and robots that can reason, learn and recognize emotion. They might finally realize a dream that is as old as the golem of Jewish folklore and as current as blockbuster science fiction: robots that understand human emotions, and that can adapt to new environments and unexpected situations.

According to Stan Franklin, a Professor in the Cognitive Computing Research Group at the University of Memphis in Tennessee, USA—who describes himself as a mathematician turned computer scientist turning cognitive scientist—these types of cognitive computing project usually have aspects of both biology and engineering. To replicate the inner workings of the brain, scientists first need to understand how the brain processes information, creates emotions and achieves

cognition—questions that are still far from being answered. Computer scientists then use this information—or as much of it as is available—to create algorithms that emulate cognition. Once the many cognitive and engineering problems are solved, it should be possible to programme computers that think, act and feel like humans.

Martin McGinnity, Professor of Intelligent Systems Engineering at the University of Ulster in Northern Ireland, has noted that a computer that could take cues from its surroundings and respond to them would have myriad more applications than can be wrested from today’s algorithm-based boxes. His team of biologists, neuroscientists, engineers and computer scientists are trying to model human perception as part of the Sensemaker project. This multi-sensory, adaptable system is inspired by the biology of human perception systems and compiles sensory information into a representation of the environment. Yet, the project also feeds back into biology, as it aims to identify the biological principles of perception and describe them mathematically for use in intelligent machines. Such computers could one day read people’s faces for cues about their emotional state; for example, an autistic person could wear a specialized emotion-sensing device to decipher friends’ and co-workers’ non-verbal messages.

Javier Movellan, head of the Machine Perception Laboratory at the University of California, San Diego, USA, builds robots with a more simple purpose in life: to be touched and smiled at by small children as often as possible. But the aims of his research team are actually far deeper than simple ‘playtime’. Movellan has found that constant human–robot interaction helps the robots to learn how to teach small children and how to interact with humans in general.

The programmes that enable the robots to register emotions—and therefore learn the best way to make children smile—are based on machine-learning algorithms that help computers get better at whatever they are programmed to do. Similar to the neurons in the human brain, the computer programmes that form the ‘brain’ of Movellan’s social robots contain small units of information and code. These units take in and process data, and then produce a signal that affects the next analysis units down the line. In this way, the artificial neural network extracts patterns and finds rules in the data it receives and changes its structure based on the information flow. The changing structure, which slowly becomes more advanced, simulates human learning.

As Movellan explained, for his robots to be able to interact with humans they must first be able to recognize a human face; all further interactions logically flow from that initial recognition. “Rather than figuring out by hand how to solve problems like finding faces, we take many examples of faces and no faces and put them in the computer to learn the difference,” he said.

To replicate the inner workings of the brain, scientists first need to understand how the brain processes information, creates emotions and achieves cognition...

The researchers started with the popular smiley face from the 1960s and slowly worked their way up from there (Susskind *et al*, 2007). They soon connected the growing abilities of the software to motors and actuators, which were installed in



© Corbis

baby robotic dolls. “One of our experiments with the baby robots was that every once in a while we would touch and play with [them] and that was enough for them to learn how humans look like,” Movellan said. To the team’s surprise, the robots quickly learned to pick out human faces so well that they could separate the human from the inhuman in comic books and other formats.

While Movellan’s team were busy teaching their algorithms, they faced a second challenge: finding potential applications. “We know very little about social interaction from a computer’s point of view,” Movellan said. “So we thought, let’s create a simple robot and throw ourselves into field conditions and progressively learn to create better robots.”

An opportunity soon presented itself when the team decided to use the dolls to help children learn. The emotion-reading robots are now in their third year of residence at a southern Californian daycare centre where they help two-year-old children to literally tell the difference between apples and oranges, and, more generally, help the toddlers learn to dis-

criminate between words. The experiment is ongoing, but so far it has been successful beyond even Movellan’s expectations. “The parents say [that] all [of a] sudden their kids are starting to point out apples and say ‘apples’, which is what the robot was helping them to do,” he said.

In addition to their programmed desire to be played with, touched and smiled at as much as possible, the robots are also programmed to attract a child’s attention in order to ask that child a question—for example, ‘where is the apple?’ The self-learning software enables the robot to adapt its behaviour to maximize its goals. Along the way, the robot learns to make sense of its environment. “In the beginning, they learn: ‘if no one [is] around maybe I should play a song kids like so they come around,’” Movellan said. “But then they learn: ‘if I play the same song many times, that’s not good because I lose my audience.’”

Once the robot dolls were exposed to the real world at the daycare centre, the scientists were also able to identify other problems in interacting with humans; for example, the robots needed to move their heads the same way that humans do. “If it moves at the right speed it’s looking at

something; but if it’s not at the right speed that magic disappears,” Movellan said. “In our social life, these type of signals are powerful and we don’t notice them until they happen.”

For Beth Crane, a doctoral student at the University of Michigan (Ann Arbor, MI, USA), this does not come as a surprise. “It’s a challenging problem because the head can move in so many ways; it’s unrestricted,” she said. Crane has combined her interest in kinematics—the movement of objects—with computer programming, and is now investigating how body movements express emotion in everyday life.

When people go for a job interview or on a first date, Crane explained, they hold their body in such a way as to inadvertently—or advertently—communicate what they feel: usually ill-at-ease. At the same time, they are trying to read the interviewer or date’s body language, for example, whether he or she is dismissive or interested. Body movements are a part of our social environment, but it is so natural we usually do not think about it. “We use our bodies to communicate intention,” Crane said.

Software that can read human emotions—and that humans can respond to

emotionally—must therefore be able to both decipher body movements and communicate emotions, either through robot bodies or rendered bodies on a screen. Crane uses motion-capture technology to analyse movements of students in states of anger, happiness and sadness, and confirms by a third-party whether the students are really expressing those emotional states in the recordings. Working backwards from the motion-capture images, Crane creates numerical programmes that depict the way the body moves in various emotional states.

The emotion-reading robots [...] help two-year-old children to literally tell the difference between apples and oranges...

The same challenge is driving research at the University of Pennsylvania's Center for Human Modelling and Simulation (Philadelphia, PA, USA). Norman Badler's team approaches the problem from the computing side by using computer graphics, inverse kinematics and detailed human models to correctly express human emotions in computer-generated animations. The general aim behind this modelling of human movements and emotions is to create more realistic avatars, for example, in virtual-reality learning environments and instructional videos, as well as animated avatars in virtual environments or better crowd representation in digitally animated movies.

While researchers such as Movellan or Crane use various methods to measure, emulate and express human emotions and behaviour, Franklin and his colleagues at the Cognitive Computer Research Group are working on a cognitive link between perception of the environment and the expression of emotions or movements; in short, a model of human cognition. They have developed basic software that represents what they call the cognitive atom—the basic unit on which higher-level cognitive processes are built—to instil larger software packets with emotions and the ability to deliberate in a human manner.

The architecture that drives this Learning Intelligent Distribution Agent (LIDA) draws from cognitive psychology and cognitive neuroscience; every autonomous agent—whether human, animal or

artificial mind—operates by means of a continued sense–cognize–act cycle. Franklin explained the idea behind LIDA: “[y]ou’re never thinking about thinking, you just do it [...] If I walk around a couch rather than bump into it, I’m using my consciousness to tell me where the couch is and how to avoid it. I don’t need to consciously think about how to avoid it.”

LIDA is an artificial mind based on a series of digital sense–cognize–act atoms that build onto one another. Most of LIDA's functions are unconscious—it does not need to map out every thought and movement in order to experience them. However, the software is also able to make decisions consciously, as humans do. By way of example, Franklin explained, “I’m thirsty, do I want water, orange juice or a beer? I decide that consciously.”

LIDA is also programmed with rudimentary emotion that help to motivate its actions and allow it to learn from environmental cues to which it can respond. As with humans, LIDA learns from emotions that result from actions. “If I’m thirsty, I get something to drink. If I’m angry I hit someone,” Franklin said. “That’s the way motivations are implemented in humans and in animals and they’re part of LIDA.”

The researchers now plan to implement LIDA in cognitive robots that should learn how to act and move autonomously. Franklin envisions an automatic tug—a tractor-like vehicle that overnight shippers could use to pull containers from the sorting area to planes or ships and back. Today, humans drive these tugs; in the future, such vehicles might operate without human intervention, simply driven by autonomous, self-learning software. But this is not just of academic interest: the US Defense Advanced Research Project Agency (Arlington, VA, USA) for example, organizes the annual Grand Challenge and Urban Challenge Race for vehicles that can autonomously find their way in an unknown environment. Obviously, the military would find great use in such vehicles.

Another use for LIDA already exists—one in which it not only learns and develops artificial feelings or emotions, but also recognizes emotions in humans and reacts to them. The web-based AutoTutor system (www.autotutor.org), based on LIDA, engages in a dialogue with the student. The animated agent speaks with intonation, facial expressions and gestures that closely mimic human movement. But it also

models its responses on the basis of the student's answer. “AutoTutor becomes aware of the emotions of the students and adapts its responses to whether the student is in the flow or frustrated or confused or what,” Franklin said. The researchers showed that some students could learn just as well with AutoTutor as they did by reading about the subject. Studies about how AutoTutor functions as compared with human tutors are ongoing (VanLehn *et al*, 2007).

However, although such applications are interesting, it still begs the question of what cognitive computing can add to the unemotional computing that has been used successfully since the dawn of the mainframe. Not surprisingly, cognitive-computing researchers face a classic chicken-and-egg conundrum: find applications for machine intelligence before starting the work, or get to work and trust that the applications will follow.

“AutoTutor becomes aware of the emotions of the students and adapts its responses to whether the student is in the flow or frustrated or confused...”

With the field in its infancy, Movellan believes that it is hard to conceive of all possible applications. Certainly, no one can foresee the full potential of cognitive computing and its applications, just when 25 years ago the first personal computers became available. “People ask me, ‘what will you use this for?’ That’s the same question my friend asked me in 1982 when I bought my first personal computer,” Movellan said. “And now the question is, ‘why wouldn’t you have one?’” The cognitive computing field is just taking off, Franklin commented; check back in 15 years.

REFERENCES

- Susskind JM, Littlewort G, Bartlett MS, Movellan J, Anderson AK (2007) Human and computer recognition of facial expressions of emotion. *Neuropsychologia* **45**: 152–162
- VanLehn K, Graesser AC, Jackson GT, Jordan P, Olney A, Rose C (2007) When are tutorial dialogues more effective than reading? *Cog Sci* **31**: 3–62

Jean Thilmann

doi:10.1038/sj.embor.7401106