Can robots have emotions?

Dylan Evans

Science fiction is full of machines that have feelings. In *2001: A Space Odyssey*, the onboard computer turns against the crew of the spaceship *Discovery 1, and* utters cries of pain and fear when his circuits are finally taken apart. In *Blade Runner,* a humanoid robot is distressed to learn that her memories are not real, but have been implanted in her silicon brain by her programmer. In *Bicentennial Man,* Robin Williams plays the part of a robot who redesigns his own circuitry so that he can experience the full range of human feelings.

These stories achieve their effect in part because the capacity for emotion is often considered to be one of the main differences between humans and machines. This is certainly true of the machines we know today. The responses we receive from computers are rather dry affairs, such as 'System error 1378'. People sometimes get angry with their computers and shout at them as if they had emotions, but the computers take no notice. They neither feel their own feelings, nor recognise yours.

The gap between science fiction and science fact appears vast, but some researchers in artificial intelligence now believe it is only a question of time before it is bridged. The new field of **affective computing** has already made some progress in building primitive emotional machines, and every month brings new advances. However, some critics argue that a machine could never come to have real emotions like ours. At best, they claim, clever programming might allow it to *simulate* human emotions, but these would just be clever fakes. Who is right? To answer this question, we need to say what emotions really *are.*

What are emotions?

In humans and other animals, we tend to call behaviour emotional when we observe certain facial and vocal expressions like smiling or snarling, and when we see certain physiological changes such as hair standing on end or sweating. Since most computers do not yet possess faces or bodies, they cannot manifest this behaviour. However, in recent years computer scientists have been developing a range of 'animated agent faces', programmes that generate images of humanlike faces on the computer's visual display unit. These images can be manipulated to form convincing emotional expressions.

Others have taken things further by building three-dimensional synthetic heads. Cynthia Breazeal and colleagues at the Massachusetts Institute of Technology (MIT) have constructed a robot called 'Kismet' with moveable eyelids, eyes and lips. The range of emotional expressions available to Kismet is limited, but they are convincing enough to generate sympathy among the humans who interact with him. Breazeal invites human parents to play with Kismet on a daily basis. When left alone, Kismet looks sad, but when it detects a human face it smiles, inviting attention. If the carer moves too fast, a look of fear warns that something is wrong. Human parents who

play with Kismet cannot help but respond sympathetically to these simple forms of emotional behaviour.

Does Kismet have emotions, then? It certainly exhibits some emotional behaviour, so if we define emotions in behavioural terms, we must admit that Kismet has some emotional capacity. Kismet does not display the full range of emotional behaviour we observe in humans, but the capacity for emotion is not an all-or-nothing thing. Chimpanzees do not display the full range of human emotion, but they clearly have some emotions. Dogs and cats have less emotional resemblance to us, and those doting pet-owners who ascribe the full range of human emotions to their domestic animals are surely guilty of anthropomorphism, but to deny they had any emotions at all would surely be to commit the opposite, and equally egregious, error of anthropocentrism. There is a whole spectrum of emotional capacities, ranging from the very simple to the very complex. Perhaps Kismet's limited capacity for emotion puts him somewhere near the simple end of the spectrum, but even this is a significant advance over the computers that currently sit on our desks, which by most definitions are devoid of any emotion whatsoever.

As affective computing progresses, we may be able to build machines with more and more complex emotional capacities. Kismet does not yet have a voice, but in the future Breazeal plans to give him a vocal system which might convey auditory signals of emotion. Today's speech synthesisers speak in an unemotional monotone. In the future, computer scientists should be able to make them sound much more human by modulating nonlinguistic aspects of vocalisation like speed, pitch and volume.

Facial expression and vocal intonation are not the only forms of emotional behaviour. We also infer emotions from actions. When, for example, we see an animal stop abruptly in its tracks, turn round, and run away, we infer that it is afraid, even though we may not see the object of its fear. For computers to exhibit this kind of emotional behaviour, they will have to be able to move around. In the jargon of artificial intelligence, they will have to be '**mobots**' (mobile robots).

In my lab at the University of the West of England, there are dozens of mobots, most of which are very simple. Some, for example, are only the size of a shoe, and all they can do is find their way around a piece of the floor without bumping into anything. Sensors allow them to detect obstacles such as walls and other mobots. Despite the simplicity of this mechanism, their behaviour can seem eerily human. When an obstacle is detected, the mobots stop dead in their tracks, turn around, and head off quickly in the other direction. To anybody watching, the impression that the mobot is afraid of collisions is irresistible.

Are these mobots *really* afraid? Or are the spectators, including me, guilty of anthropomorphism? People once asked the same question about animals. Descartes, for example, claimed that animals did not really have feelings like us because they were just complex machines, without a soul. When they screamed in apparent pain,

they were just following the dictates of their inner mechanism. Now that we know that the pain mechanism in humans is not much different from that of other animals, the Cartesian distinction between sentient humans and 'machine-like' animals does not make much sense. In the same way, as we come to build machines more and more like us, the question about whether or not the machines have 'real' emotions or just 'fake' ones will become less meaningful. The current resistance to attributing emotions to machines is simply due to the fact that even the most advanced machines today are still very primitive.

Some experts estimate that we will be able to build machines with complex emotions like ours within fifty years. But is this a good idea? What is the point of building emotional machines? Won't emotions just get in the way of good computing, or even worse, cause computers to turn against us, as they so often do in science fiction?

Why give computers emotions?

Giving computers emotions could be very useful for a whole variety of reasons. For a start, it would be much easier and more enjoyable to interact with an emotional computer than with today's unemotional machines. Imagine if your computer could recognise what emotional state you were in each time you sat down to use it, perhaps by scanning your facial expression. You arrive at work one Monday morning, and the computer detects that you are in a bad mood. Rather than simply asking you for your password, as computers do today, the emotionally-aware desktop PC might tell you a joke, or suggest that you read a particularly nice email first. Perhaps it has learnt from previous such mornings that you resent such attempts to cheer you up. In this case, it might ignore you until you had calmed down or had a coffee. It might be much more productive to work with a computer that was emotionally intelligent in this way than with today's dumb machines.

This is not just a flight of fancy. Computers are already capable of recognising some emotions. Ifran Essa and Alex Pentland, two American computer scientists, have designed a program that enables a computer to recognise facial expressions of six basic emotions. When volunteers pretended to feel one of these emotions, the computer recognised the emotion correctly ninety-eight per cent of the time. This is even better than the accuracy rate achieved by most humans on the same task! If computers are already better than us at recognising some emotions, it is surely not long before they will acquire similarly advanced capacities for expressing emotions, and perhaps even for feeling them. In the future, it may be humans who are seen by computers as emotionally illiterate, not vice versa.

What other applications might there be for emotional computers other than providing emotionally intelligent interfaces for desktop PCs? Rosalind Picard, a computer scientist at the MIT Media Laboratory in Boston, has proposed dozens of possible uses, including the following:

- Artificial interviewers that train you how to do well in job interviews by giving you feedback on your body language
- Affective voice synthesisers that allow people with speech problems not just to speak, but to speak in genuinely emotional ways
- Frustration monitors that allow manufacturers to evaluate how easy their products are to use
- Wearable computers ('intelligent clothing') that give you feedback on your emotional state so that you can tell when you are getting stressed and need a break

All of these potential applications for emotional machines are resolutely utilitarian, but I think that most emotional machines in the future will be built not for any practical purpose, but purely for entertainment. If you want to envision the future of affective computing, don't think spacecraft and intelligent clothing – think toys and videogames.

Many videogames already use simple learning algorithms to control non-player characters, such as monsters and baddies. In *Tomb Raider*, for example, the enemies faced by Lara Croft need only a few shots before they cotton on to your shooting style. If you are lying in wait for a dinosaur, it might remain in the shadows, tempting you to come out and take a pot shot so that it can attack you more easily. These are relatively simple programs, but the constant demand for better games means that the software is continually improving. It might well be that the first genuinely emotional computers are games consoles rather than spacecraft.

Other entertainment software with proto-emotional capacities is also available in the form of the virtual pets who live in personal computers. Many kids now keep dogs and cats as screen-pets, and more recently a virtual baby has been launched. A program called the Sims lets you design your own people, but they soon take on a life of their own, which can be fascinating to watch. The Sims are eerily human in their range of emotional behaviour. They get angry, become depressed, and even fall in love.

All these creatures are virtual – they live inside your computer, and their only 'body' is a picture on a screen. However, the first computerised creatures with real bodies are also now coming onto the toy market, and they too have proto-emotional capacities. First came little furry robots called 'Furbies', that fall asleep when tired, and make plaintiff cries when neglected for too long. Now there are also robotic dogs and cats that run around your living room without ever making a mess. There is even a baby doll with a silicon brain and a latex face that screws up into an expression of distress when it needs feeding. As with Kismet, people respond to these artificial life forms with natural sympathy. Their minds are not filled with ponderous doubts about whether these emotions are 'real' or not. They simply enjoy playing with them, as they would with a real kitten or baby.

The gap between science fiction and science fact is closing. Today's computers and

robots still have a long way to go before they acquire a full human range of emotions, but they have already made some progress. In order to make further progress, engineers and computer scientists will have to join forces with psychologists. Increasing numbers of psychology students are opting to study robotics and artificial intelligence at university. The future lies in their hands.

Word count: 2083

Acknowledgements

This article is a revised and edited version of chapter five in *Emotion: The Science of Sentiment*, by Dylan Evans (Oxford University Press, 2001). Thanks are due to Oxford University Press for granting permission to reproduce material from this book.

References

Breazeal, C. (2002) *Designing Sociable Robots.* MIT Press.

Essa, I. and A. Pentland (1997). "Coding, analysis, interpretation and recognition of facial expressions." *IEEE Transactions on Pattern Analysis and Machine Intelligence* **19**(7): 757-763.

Evans, D. (2001) *Emotion: The Science of Sentiment.* Oxford University Press.

Grand, S. (2004) *Growing up with Lucy: How to build an android in twenty easy steps*. Weidenfeld & Nicholson.

Picard, R. W. (1997). *Affective Computing*. MIT Press.

Note on the author: Dylan Evans teaches robotics at the University of the West of England, Bristol. He has written several popular science books, and writes regularly for *Guardian* and the *Evening Standard.* In 2001 he was voted one of the twenty best young writers in Britain by the *Independent on Sunday* , and was recently described by the *Guardian* as 'Alain de Botton in a lab coat'. See www.dylan.org.uk.

Key concepts:

- affective computing
- mobots
- intelligent clothing
- artificial intelligence
- artificial life