12.

Fixed points in a changing world

Peter Robinson

University of Cambridge

Abstract

The importance of emotional expression as part of human communication has been understood since the seventeenth century, and has been explored scientifically since Charles Darwin and others in the nineteenth century. Recent advances in Psychology have greatly improved our understanding of the role of affect in communication, perception, decision-making, attention and memory. At the same time, advances in technology mean that it is becoming possible for machines to sense, analyze and express emotions. We can now consider how these advances relate to each other and how they illuminate human nature.

Keywords: Affective Computing, Robots, Relationships.

Introduction

Humans are made for relationships – a relationship with God and relationships with each other – and emotion plays a crucial role in those relationships. The ability to display and recognize emotions is an important aspect of human communication. We monitor each other's facial expressions, vocal nuances and body posture and gestures, and make inferences from them about each other's mental states. Computers are gaining the ability to display and recognize human emotions, and it is tempting to think of the machines as actually having emotional intelligence. But they are only synthetic emotions and we must avoid the trap of thinking that humans can have meaningful relationships with machines.

This distinction concerning synthetic emotion is significant because human beings are social animals. Our interactions with other people are informed by the inferences that we make about their emotions from facial expressions, vocal expression, and body posture and gestures. This understanding of mental states shapes the decisions that we make, governs how we communicate with others, and affects our performance. The ability to attribute mental states to others from their behavior, and to use that knowledge to guide our own actions and predict those of others is described by psychologists as 'theory of mind' or 'mind-reading'. It has recently gained attention with the growing number of people with Autism Spectrum Conditions, who have difficulties mind-reading.

People express these social signals even when we are interacting with machines, but computer interfaces currently ignore them. Computers lack emotional intelligence. Recent advances in psychology have greatly improved our understanding of the role of affect in communication, perception, decision-making, attention, and memory. At the same time, advances in technology mean that it is becoming possible for machines to sense, analyze and express emotions. "Affective computing" (Picard, 1997) explores the relationship between these advances and is bringing them together to endow computers with emotional intelligence.

Mind-reading computer systems have been developed that infer mental states such as enjoyment, agreement, interest, and confusion from facial expressions in real-time by using a combination of computer vision, machine learning, and software engineering. The applications encompass all aspects of human-computer interaction (Robinson et al., 2011). On-line teaching systems can monitor a student and adjust the pace and content of a lesson as it detects interest or boredom, understanding or confusion. Telematic systems in cars can monitor the driver's cognitive load and suppress interruptions from the vehicle when the driver is overloaded. Medical applications include diagnosing depression or sensing pain. Teleconference systems can use the information to break down the artificial barriers presented by a screen.

Commercial applications of the technology are beginning to arrive on the market. Systems have been deployed on a large scale to evaluate video content by measuring audience engagement with media presentations and tracking their responses to brand identities. The increasing sensor capabilities and processing power on smartphones have allowed the technology to be implemented on mobile platforms, enabling a host of new digital experiences from games that adapt to your emotions to wellness apps that monitor your mood (Marchi et al., 2018).

At the same time there has been a growing interest in computer systems that express emotions. These take the form both of android robots with expressive features and of text-based systems that empathize with their users. These, combined with a growing use of artificial intelligence systems based on applying machine learning to large volumes of data, have led to speculation about sentient machines forming relationships with humans. The speculation is mainly in the form of fictional narratives, from Karel Čapek's *Rossum's Universal Robots* in 1920, through Michael Crichton's *Westworld* in 1973 to the current television series *Humans*. These speculate that the growing power of machine intelligence will somehow make computers sentient, but this is used as a vehicle to explore human relationships rather than as an exposition of technology.

Mental States

Charles Darwin (1872) published *The Expression of the Emotions in Man and Animals* in 1872, exploring the role of emotional expression in communication between humans. Over a century later, Rosalind Picard at MIT observed that effective communication between people and computers also requires emotional intelligence; computers must have the ability to recognize and express emotions. The study of affective computing has blossomed subsequently.

Figure 12.1: Mental States identified in *The Expression of the Emotions in Man and Animals*

		
Boredom	Confusion	Fear
		
Happiness	Sadness	Surprise

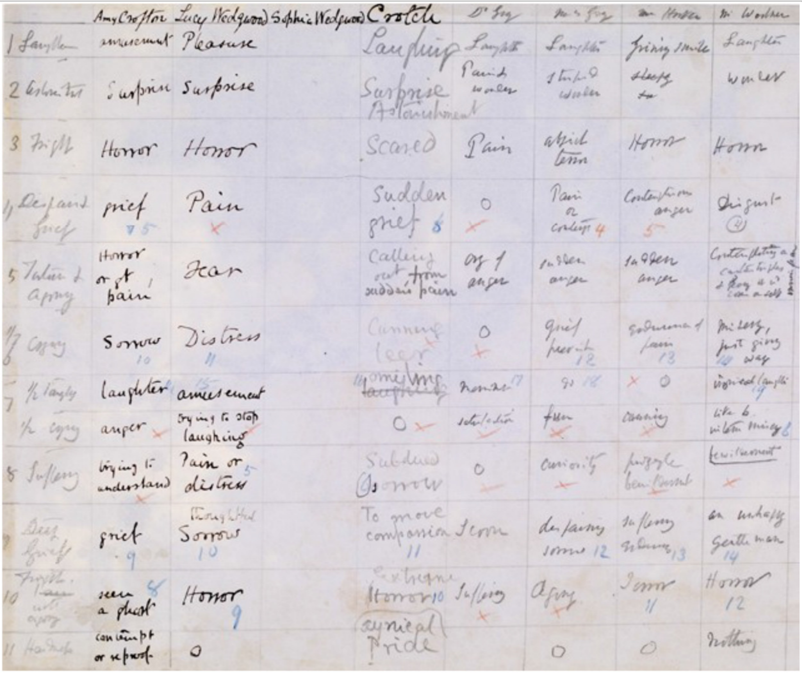

Darwin’s spreadsheet

Darwin was interested in the universality of emotional expressions, which could give an evolutionary advantage to humans. He investigated this using photographs from the French scientist Guillaume Duchenne de Boulogne who had been considering the stimulation of facial muscles using electric shocks. Darwin invited guests at his house to examine the photographs and say what emotion they saw in them. The results were entered in a spreadsheet, with rows for each picture and columns for each assessor. Their agreement convinced Darwin that the expressions were, indeed, universal. A recent study using the same photographs and web-based crowd-sourcing showed that the photographs elicit the same labels 140 years later (Mahmoud et al., 2012). Human inference of basic emotions from facial expressions remains universal and unchanging.

Darwin considered seven categories of emotion in his work. A century later, Paul Ekman at the University of California refined this into a classification of six basic emotions – anger, disgust, fear, joy, sadness and surprise (Ekman et al. 1972). The six basic emotions and Ekman's Facial Action Coding System (FACS) have been widely used in the study of emotions over the past 35 years (see Ekman and Friesen, 1978), and particularly for work on affective computing in the past 15 years. However, they are not particularly representative of people's everyday experiences.

Recent work by Simon Baron-Cohen, who directs the Autism Research Centre at the University of Cambridge, has led to a new taxonomy of human emotions based on a linguistic analysis (Baron-Cohen et al., 2002). 412 distinct emotion concepts were identified and grouped into 24 disjoint categories. These broader categories include Ekman's six basic emotions and a further 18 groups that cover complex mental states reflecting cognitive activity – conditions such bored, interested, sure, unsure, thinking and so on. These require a few seconds of continuous observation to be recognized by humans, rather than the single image that suffices for basic emotions. However, they are more representative of people's everyday emotions.

James Russell at the University of British Columbia took a different approach by deriving a continuous, dimensional classification in his 'Circumplex' model of affect (Russell, 1980). This was formulated in the light of an experiment in which participants arranged 28 emotion words around a circle, with similar affects located close to each other and inverses on opposite sides of the circle. Principal Component Analysis was then used to identify various dimensions in the data. The first two components accounted for 46% of the total variance, and the next three only an additional 13%.

These two components are usually referred to as 'valence' (running from negative to positive) and 'arousal' (running from passive to active). The further axes have been given names like 'intensity', 'expectancy' and 'tendency' (inward or outward). This has led to a popular belief that emotions can be measured precisely by coordinates in a suitably high-dimensional space. Unfortunately, this is not true for machines any more than it is for humans, and any computation involving emotions must be designed to handle ambiguity and uncertainty. A charming example is the 1909 promotional material for the actor Florence Lawrence, Hollywood's 'Biograph Girl' (Blum, 1953). Facial expressions were crucial in the era of silent films, and this shows her ability to portray a wide variety of emotions, although 'Piety' and 'Sadness' are indistinguishable. This also illustrates the importance of dynamic information rather than just still images in understanding faces (el Kaliouby et al., 2003).

Recognizing Emotions

Although Darwin (1872) concentrated on facial features to convey emotions, he also mentioned vocal sounds, other sounds, body posture and gesture, and physiological responses as further indications of emotion. All of these channels have been considered as ways of automatic monitoring emotion in humans, although these sensors used for some are more invasive than for others. Wiring a person with electrodes to measure their heart rate, breathing, and skin conductivity is likely to provoke unwanted emotions! Signals that can be monitored non-invasively using cameras and microphones are much more suitable.

Facial Expressions

People routinely express their mental states through their facial expressions and this is one of the clearest channels for communication. Inference from facial expressions has been studied using a variety of techniques – rule-based classifiers, neural networks, support vector machines, and Bayesian classifiers – but mostly restricted to the six basic emotions. Recognizing complex, cognitive mental states is more difficult, but probably more useful as part of general interaction with computer systems. It is now possible to build a fully automatic system for recognizing emotions that requires no human intervention and which operates in real-time on commodity hardware.

Figure 12.2: Facial affect inference.

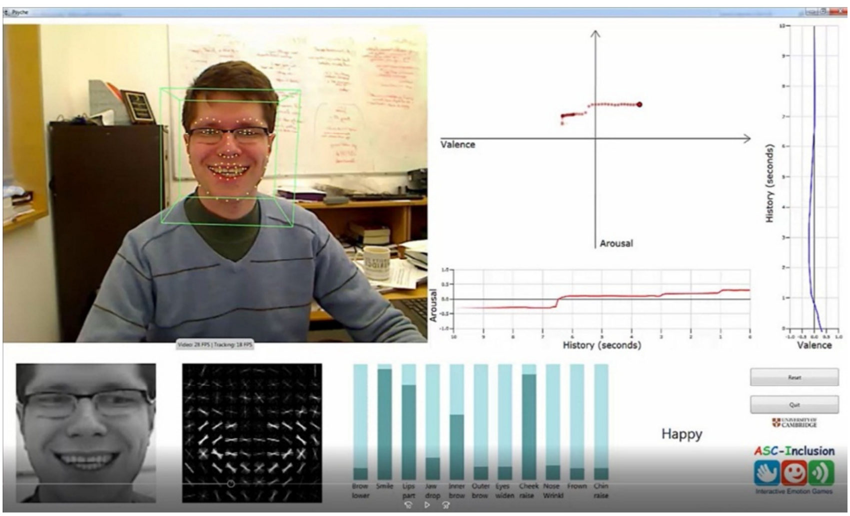

The webcam image is shown at the top left with some tracking information superimposed in green. Various stages of processing are shown across the bottom, with histograms for strengths of action units and a final classification. The graphs at the top right show continuous measures of valence and arousal.

A great deal of data is needed to determine the timing characteristics of people’s expressions and to train the statistical classifiers used in the inference system. Baron-Cohen’s Mind Reading DVD (assembled to teach children with high-functioning autism spectrum conditions to recognize emotions) proved ideal for this purpose. An evaluation considered six conditions drawn from five of the 24 emotion groups and including 29 of the underlying mental state concepts, and chosen to be particularly relevant for human-computer interaction (el Kaliouby Robinson, 2004). For a mean false positive rate of 5%, the overall accuracy of the system is 77%. The system also generalizes well to faces not included in the training data (el Kaliouby Robinson, 2005). An alternative system estimates continuous measures of valence and arousal.

Non-Verbal Aspects of Speech

The voice provides another significant channel for the expression of emotions. Features such as the pitch, energy and tempo can reveal a lot about the mood of the speaker. There are no characteristic features that indicate particular mental states but it is possible to distinguish between using two emotions using a small number of features, with a different set of features may be required to distinguish those emotions from others. The most successful approach is to calculate a large collection of about 170 features for each utterance. A training phase uses data mining to identify the features that separate each pair of emotional conditions. The operational phase then uses these pair-wise comparisons as preferences in a voting scheme to give an overall ranking (Sobol Shikler Robinson, 2010).Evaluation separated nine conditions with an accuracy of 70%, increasing to 83% if multiple winners were considered. The approach also generalizes well to speakers other than those used for training, and even to other languages.

Body Posture and Gesture

The third natural channel for expression of emotions includes body posture and gesture (Bernhardt and Robinson, 2007). However, characteristics indicating what movement is being considered and which person is doing it must be discounted before it is possible to analyze how it is being done (Bernhardt and Robinson, 2009). Movement involves an individual bias, so the analysis is harder than for facial expressions or voice.

The solution is to break complex motions down into a system of isolated elements whose dynamic cues can be used to distinguish affects. This is similar to the process of breaking continuous speech into phonemes. As with affective analysis of speech, pair-wise comparisons are used on individual motion segments, and each segment is classified using a majority vote. A complete motion is then classified by a majority vote of its component segments. The method was tested on a corpus of about 1200 motion samples, representing roughly equal numbers of four expressions of four different actions. The average recognition rate of 81% is comparable to the rates achieved by human observers of similar data.

Expressing Emotions

Until recently, robots have been separated into two quite separate categories – industrial robots used in manufacturing that are powerful but need to be isolated behind safety barriers, and domestic robots that meander round the home but are too weak to do much more than clean the floor and serve drinks. However, service robots with sufficient strength to be harmful as well as useful may soon be deployed in domestic environments. A typical application might be care of the elderly at home where the robot would assist a health care professional with tasks that require physical strength. The robot would need to be sensitive to the unspoken mental states of both the patient and the carer, and must also reassure them through its own expression.

Humans routinely convey empathetic responses through involuntary facial mimicry, and this extends to human-robot interaction. An experiment showed that conversation between a participant and a robot is enhanced when the robot mimics the subject rather than moves randomly (Riek et al., 2010). This raises questions about the degree of human-likeness required in the appearance of robots that interact with humans. A further experiment investigated participants' empathy for robots shown in film clips, and the responses were directly correlated with their resemblance to humans (Riek et al., 2009).

Figure 12.3: ***Charles***

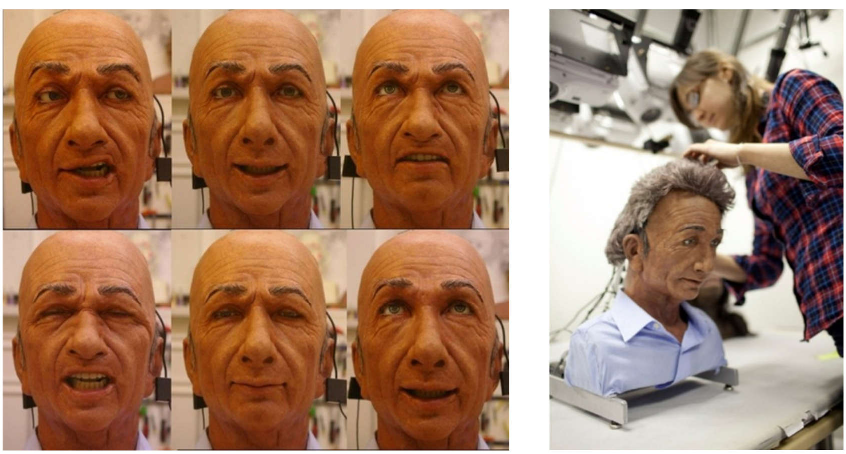

Charles, a robotic head made by David Hanson with 28 motors replicating muscles in the human face.

Applications

Autism Spectrum Conditions (ASCs) are neurodevelopmental conditions characterized by social communication difficulties and restricted and repetitive behavior patterns. The European ASC-Inclusion project worked to create and evaluate the effectiveness of an internet-based game platform, intended for children with ASCs and their carers (Schuller et al., 2015). The platform combines several state-of-the-art technologies in one comprehensive virtual world providing training through games, and including feedback from analysis of the player's gestures, facial and vocal expressions using a standard webcam and microphone. The game also includes text communication with peers and smart agents, animation, video and audio clips (Marchi et al, 2018).

Figure 12.4: The ASC-Inclusion training system.

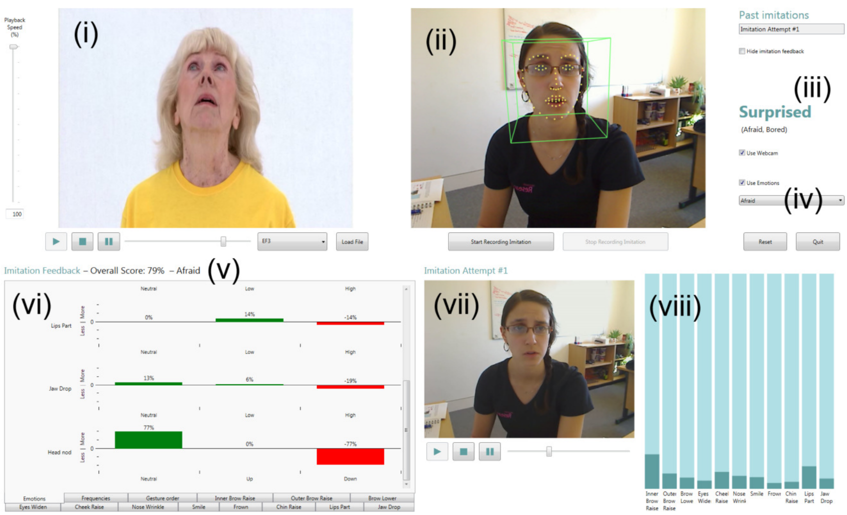

The images show (i) an action to be imitated and (ii) the student's attempt, together with diagnostic information suggesting changes required to improve the expression.

Another application of affective inference is monitoring cognitive load in command-and-control systems. Driving a car provides a good model for this. Driving is becoming increasingly difficult with increasing traffic densities combined with distractions from in-car technologies such as mobile phones and satellite navigation systems. Drivers coping with heavy traffic in an unfamiliar city while late for a meeting are not helped by a navigation system that instructs them to make a U-turn. An emotionally aware car would help by suppressing phone calls, turning the radio off, and even allowing drivers to proceed in the wrong direction until they had recovered their composure (Wright et al., 2017).

However, it is difficult to construct repeatable experiments using real cars on real roads. A common approach is to use simulation which allows controlled experiments, but fails to engage participants. An alternative possibility is the use of remotely controlled vehicles. Participants are located in the laboratory where they can be monitored easily, while controlling a real vehicle that is undertaking a task in a real environment. We have found that remotely controlled helicopters engender a particularly strong sense of emotional investment (Davies and Robinson, 2011).

In all of this, it is important to remember that inferring someone's mental state is not a precise science. People routinely misread each other's social signals and it would be foolish to expect computers to be any more accurate. Expressions of emotions are inherently ambiguous and using this sort of information in automatic systems requires careful consideration of human factors as well as intelligent use of probabilistic computing (Robinson and Baltrušaitis, 2015).

In particular, an affective inference system should not be regarded in the same way as a piece of precision measuring equipment. It seems unlikely that it will ever be possible to point a camera at somebody and read their emotions. A more practical approach is to formulate a set of perhaps half a dozen conditions to be distinguished in a particular application. For example, it might be useful to know if a car driver is comfortable, pleased, bored, drowsy, concentrating, confused, upset or, indeed, none of these. Each of these conditions would be populated with 5 to 10 of the emotion concepts in Baron-Cohen's taxonomy, and the machine learning systems trained to distinguish them.

The resulting analysis would only operate across a small subset of the entire gamut of human emotions, and statistical measures could be calculated to indicate the confidence with which the conditions could be separated. The same statistics could be used to attach "signal strength" indications to the inferences, making the ambiguity clear to any other systems that relied on them. Many modern computing systems require this sort of careful engineering to handle uncertainty in a principled way. The steadily increasing power of computing hardware, combined with reductions in size and power consumption, means that these sort of systems will soon be usable on portable equipment such as smartphones. This will open up many exciting applications from games to health care.

Public Perceptions

These technological advances are based on steady progress with scientific understanding of computing techniques and human psychology. Public perceptions, on the other hand, have leaped ahead, based as much on fiction as on fact. There is a long history in fiction of men creating beings in their own image – from the golem mythology of Jewish folklore, through Mary Shelley's *Frankenstein*, Karel Čapek's *Rossum's Universal Robots*, Rotwang's robotic Maria in Fritz Lang's *Metropolis*, and Isaac Asimov's robot series with its three laws, through to modern drama such as *Westworld* and *Humans*. In all of these, the machines are depicted as resembling humans and develop sentience independently from their masters, usually turning against their creators. These depictions are curiously at odds with robots in real life.

Robots are widely deployed in industry. They undertake repetitive tasks requiring strength and accuracy in making cars, they manipulate integrated circuits with precision beyond human capability, they assist surgeons in undertaking minimally invasive surgery with delicacy and subtlety. These are large and expensive machines that work in specialized environments, have little autonomy and display no sentience.

Robots are also beginning to appear in the home. They take the form of automatic vacuum cleaners or grass mowers, together with novelty machines that serve drinks and so on. They have no great strength and pose little threat to humanity.

Robots are popular with the military. They are used for bomb disposal, aerial surveillance and to extract wounded combatants from battlefields. These mechanical systems are also being combined with weapon systems to make remotely controlled "soldiers" on the ground and in the air. The nature of warfare changes when the human participants are safely ensconced in a bunker many miles from the site of hostilities.

Vint Cerf has characterized a robot as any system that ingests information, processes it, and produces outputs that have perceptible effects (Cerf, 2013). Such robots encompass computer systems that are less visible but potentially more harmful. It includes the global telephone system including mobile handsets, it includes the satellite navigation systems on which we rely increasingly, it includes automatic trading systems in the financial markets. These have little physical presence and are certainly not humanoid, but they exercise autonomy and wield considerable power over our lives.

Combining mobile weapon systems with autonomous control raises the specter of robot armies, and it begins to look as if some of fictional nightmares might be approaching reality. But it is important to remember that these systems are not sentient and there is little likelihood of their acquiring sentience by themselves. They may be used to perpetrate evil acts, but that does not mean that they have evil intentions of their own.

Systems linking sensors and servos under autonomous control are not sentient in any real sense. To impute them with personhood is a category error. Such systems do raise ethical issues, but they are the same ethical issues as those confronting any profession: the questions of motivation and competence. Autonomous weapons raise serious ethical questions of motivation and automatic trading systems raise serious ethical questions about competence.

However, the public perception remains. The word robot conjures up an image of a humanoid machine that is sentient and probably malicious. Why do we have this fascination with robots? Does it perhaps tell us more about humans than it does about machines?

Religious Considerations

One explanation is a human desire to make machines in our own image. This is to assume the divine ability to create men in God's image without understanding what God's image means. A naïve understanding might look for a physical resemblance, and that naivety partly explains the human perception of robots in humanoid form.

A more sophisticated reading of Genesis sees the image of God being seen in relationships and in ruling. God made mankind for relationships, both with each other and with God: The Lord God said, "It is not good for the man to be alone. I will make a helper suitable for him" (Gen 2:18) and "The man and his wife heard the sound of the Lord God as he was walking in the garden in the cool of the day" (Gen 3:8).

However, the image in robots is a poor one. The physical resemblance of robots to humans is weak and there is no relationship either between one robot and another or between robots and their creators. Robots that can recognize and display emotions may simulate empathy but do not experience it. A person who recognized and displayed emotions without understanding them would be called a psychopath. The confusion in public perceptions about robots could be as troubling for humanity as psychopaths moving freely in society.

More importantly there is no spiritual aspect to the relationship. God also made mankind to rule the world. "God said, 'Let us make mankind in our image, in our likeness, so that they may rule over the fish in the sea and the birds in the sky, over the livestock and all the wild animals, and over all the creatures that move along the ground"(Gen 1:26). We do make robots to rule the world for us, but this is generally constrained by guidance that we build into them. However, there are two problems with this. Automatic systems may be used for malicious purposes, but that malice originates with the human designers of the systems, not in the machines themselves. More subtly, modern computing techniques involving machine learning allow that guidance to be increasingly imprecise, which can lead to unintended consequences. But we should not fall into the trap of thinking that the machines have become sentient. It is simply a failure of the human designers to implement the systems correctly. In other words, failures in machines to rule properly are simply reflections of professional failures by their designers, the classic human failures of motivation and competence.

Humankind's relationship with God and delegated rule were broken at the fall. We no longer enjoy the same relationship with God and our rule is imperfect. Failures by computer systems are not a sign of original sin, but are a reflection of imperfections in their human creators.

This puts the nascent 'robot rights' movement into perspective. Robots are not genuinely free agents but are merely remotely controlled by their human designers. They may simulate appreciation and expression of emotions, but it is only a simulation. They may follow our direction in undertaking tasks, but that is not the same as exercising the sort of responsibility that earns rights.

John Wyatt exposes the difference between men and machines most clearly in referring to the Nicean Creed written in AD 325 that declared Jesus to be "begotten not made" (Wyatt and Robinson, 2019). Jesus was fully human as well as fully God. The distinction between begotten and made helps us distinguish between humankind and machines. Humans are begotten but machines are made. They are entirely different.

Conclusion

The robots are coming. Hardly a week goes by without another headline warning of the dangers of our increasing use of technology. And there are, indeed, dangers. But they are not the dangers of sentient machines in humanoid form taking over the world and either turning mankind into their slaves or simply eliminating humans. The dangers of the machines are simply the very human risks of motivation and competence in their creators. These risks must be understood and managed. However, the public perceptions of robots are still one of humanoid machines, usually with malicious intent. This tells us more about human nature than it does about the machines, and is an interesting reflection of our own fallen nature. Our relationship with our creator God remains a fixed point as the world around us changes.

References

Baltrušaitis, T., L. Riek, and P. Robinson., 2010. *Synthesising expressions using facial feature point tracking: How emotion is conveyed.* Florence, Italy: ACM Workshop on Affective Interaction in Natural Environments.

Baron-Cohen, S., O. Golan, S. Wainwright, and J. Hill., 2002. *Mind reading: the interactive guide to emotions.* DVD ed. London: Jessica Kingsley.

Bernhardt, D. and P. Robinson., 2007. *Detecting affect from non-stylised body motions.* Lisbon, Portugal: Affective Computing and Intelligent Interaction (ACII), pp. 59-70.

Bernhardt, D. and P. Robinson, 2009. *Detecting emotions from connected action sequences.* Kuala Lumpur, Malaysia: UKM International Visual Informatics Conference.

Blum, D., 1953. *A pictorial history of the silent screen.* London: Spring Books.

Cerf, V., 2013. What's a robot? *Communications of the ACM* 56(1).

Costa, S., H. Lehmann, K. Dautenhahn, B. Robins, and F. Soares., 2015. Using a humanoid robot to elicit body awareness and appropriate physical interaction in children with Autism. *International Journal of Social Robotics* 7(2), 265-278.

Darwin, C., 1872. *The expression of the emotions in man and animals.* London: John Murray.

Davies, I. P. and P. Robinson., 2011. *Emotional investment in naturalistic data collection.* Memphis TN: Affective Computing and Intelligent Interaction (ACII).

Ekman, P. and W.V. Friesen., 1978. *Facial action coding system: A technique for the measurement of facial movement.* San Francisco CA: Consulting Psychologists Press.

Ekman, P., W.V. Friesen, and P. Ellsworth., 1972. *Emotion in the human face.* Newyork: Pergamon Press.

El Kaliouby, R., P. Robinson, and S. Keates., 2003. *Temporal context and the recognition of emotion from facial expression.* Crete: HCI International.

El Kaliouby, R. and P. Robinson., 2004. *Real-time inference of complex mental states from facial expressions and head gestures.* Washington DC: IEEE Computer Visions and Pattern Recognition (CVPR).

El Kaliouby, R. and P. Robinson., 2005. *Generalization of a vision based computational model of mind-reading.* Beijing, China, Affective Computing and Intelligent Interaction (ACII), pp. 582-589.

Mahmoud, M., T. Baltrušaitis, and P. Robinson., 2012. *Crowdsouring in emotion studies across time and culture.* Nara, Japan: ACM Workshop on Crowdsourcing for Multimedia.

Marchi, E. et al., 2018. The ASC-Inclusion perceptual serious gaming platform for autistic children. *IEEE Transactions on Computational Intelligence in Games.*

Picard, R. W., 1997. *Affective Computing.* Boston: MIT Press.

Riek, L. D., P. Paul, and P. Robinson., 2010. When my robot smiles at me: Enabling human-robot rapport via real-time head gesture mimicry. *Journal on Multi-modal User Interfaces,* 3(1-2), pp. 99-108.

Riek, L. D., T. Rabinowitch, B. Chakrabarti, and P. Robinson., 2009. *Empathizing with robots: Fellow feeling along the anthropomorphic spectrum.* Amsterdam: Affective Computing and Intelligent Interaction (ACII).

Riek, L. D. and P. Robinson., *Synthesising head and facial movement disorders on android robots.* Edinburgh, Scotland: ACM Symposium on Facial Analysis and Animation.

Robinson, P. and T. Baltrušaitis., 2015. *Empirical analysis of dimensional affect.* Xi'an, China, Affective Computing and Intelligent Interaction (ACII).

Robinson, P. et al., 2011. *The emotional computer.* San Francisco CA: ACM Conference on Pervasive Computing.

Russell, J. A., 1980. A circumplex model of affect. *Personality and Social Psychology,* 39(6), pp. 1161-1178.

Schuller, B. et al., 2015. *Recent developments and results of ASC-Inclusion.* Atlanta GA: ACM Intelligent User Interfaces.

Sobol Shikler, T. and P. Robinson., 2010. *IEEE Transactions on Pattern Analysis and Machine Intelligence,* 32(7), pp. 1284-1297.

Wright, J. et al., 2017. *Intelligent scheduling for in-car notifications.* Modena, Italy: IEEE Forum on Research and Technologies for Society and Industry.

Wyatt, J. and P. Robinson, 2018. *God, artificial intelligence and humans: Could robots replace humans made in the image of God?.* Cambridge: Faraday Institute.