Synthetic emotions and the *imago Dei*

Peter Robinson
Gonville & Caius College, Cambridge

The robots are coming. Turn on the television and you will see series like Humans or Westworld where human-like robots turn on their creators. Read the newspaper and you will see headlines about robots taking our jobs or holding beauty contests for people. In one case an intelligent robot that remembers and learns was even reported to have escaped from its home laboratory on more than one occasion (Wajeeha, 2016). Serious scientists such as Stephen Hawking and the astronomer royal Martin Rees are joined by industry leaders like Bill Gates and Elon Musk in talking about artificial intelligence developing so quickly that it will take over from humans (Hawking, 2014). Recent developments in the use of large language models for generative conversational systems such as ChatGPT have led to speculation that sentient machines are already here. Some even talk of this being the next step in evolution on Earth, and we humans will become as extinct as the dinosaurs as computers take over.

What is the proper Christian response to these questions? As ever, we must turn to the Bible for guidance.

In the beginning…

The Bible, of course, says nothing explicit about the philosophical, social and theological implications of automation, robots and artificial intelligence. However, it does give us a framework within which we can assess the value of human endeavours. The opening chapters of Genesis set out God's intention of what it means to be human. They can help us put all these new technologies into perspective. And the whole Bible explains God's greater purpose for humans in his creation.

Genesis is not written as a historical narrative or as a scientific textbook, so we shouldn't try to take the six days of creation too literally. But it does tell us a lot about God – the word 'God' occurs 30 times in the first 31 verses. And it tells us about God's creation, particularly the pinnacle of creation, which is mankind.

Genesis 1:26 says:

> *Then 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."*

God created people in his own image, but what is that image? The best place to look is right there in the earlier verses of Chapter 1 of Genesis. They tell us three important things about God.

First they tell us that God was there in the beginning. In the beginning God… I am not a physicist but I understand that one of Stephen Hawking's last scientific papers, published after his death, describes how the mathematics of the big bang means that there actually was a beginning to time (Hawking, 2018). Time as we understand it started at that moment of creation. God was there in the beginning.

Secondly, God is creative. God created the heavens and the earth, he created light and darkness, he brought order to the chaos, he separated land and sea, he created plants, fish, birds and wild animals. Finally, he created mankind.

Thirdly, we see God as a ruler and a judge. Through the six days of the creation, God gives orders and those orders are obeyed. Six times we read, "And it was so." God rules. God also judges. Five times we read "And God saw that it was good." And finally, "God saw all that he had made, and it was very good."

These opening verses of Genesis portray God as a creator and ruler, with mankind – you and me – the pinnacle of that creation and made in that image. In verse 28, God's first words to mankind instruct them to:

> *"Be fruitful and increase in number; fill the earth and subdue it. Rule over the fish in the sea and the birds in the sky and over every living creature that moves on the ground."*

In other words, mankind is to be creative, to be fruitful and to increase in number, and to rule over creation. We are made in the image of God to be creative and to rule.

Chapter 2 of Genesis introduces another aspect of God's image that is reflected in mankind – the need for relationships. In verse 18:

> *The Lord God said, "It is not good for the man to be alone. I will make a helper suitable for him."*

It is not good to be alone, we are made for relationships. There is a wonderful picture in the next couple of verses describing God and Adam working together as they tidy up the details of creation:

> *Now the Lord God had formed out of the ground all the wild animals and all the birds in the sky. He brought them to the man to see what he would name them; and whatever the man called each living creature, that was its name. So the man gave names to all the livestock, the birds in the sky and all the wild animals.*

We have a lovely picture of God and the man playing this naming game together. God brings the animals and birds to Adam, and Adam names them. It's a close collaborative relationship between man and his creator as Adam helps God.

But who is going to help Adam? God's answer is another person, woman. Man and woman are made for a special relationship with each other. A man leaves his father and mother and is united to his wife, and they become one flesh. People are made for relationships – relationships with God and relationships with each other, especially between man and wife.

After the first two chapters of Genesis, we see humans in the image of God, made to be creative and to rule, and made for relationships with God and with each other. Rule and relationships. It's a wonderful picture.

But then comes Chapter 3 and the fall. Mankind shows itself unfit to rule and breaks the relationship with God. That is our world today – a world of imperfect rule and broken relationships.

Where do robots and artificial intelligence fit in that imperfect, broken world?

Synthetic emotions

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 behaviour, and to use that knowledge to guide our own actions and predict those of others is known 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, analyse and express emotions. *Affective computing* 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 (el Kaliouby & Robinson, 2004). The applications encompass all aspects of human-computer interaction. 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 has allowed the technology to be ported to mobile platforms, enabling a host of new digital experiences from games that adapt to your emotions to wellness apps that monitor your mood.

Mental states

Charles Darwin published *The expression of the emotions in man and animals* in 1872 (Darwin, 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 (Picard, 1997). The study of affective computing has blossomed subsequently.

	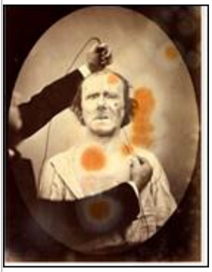		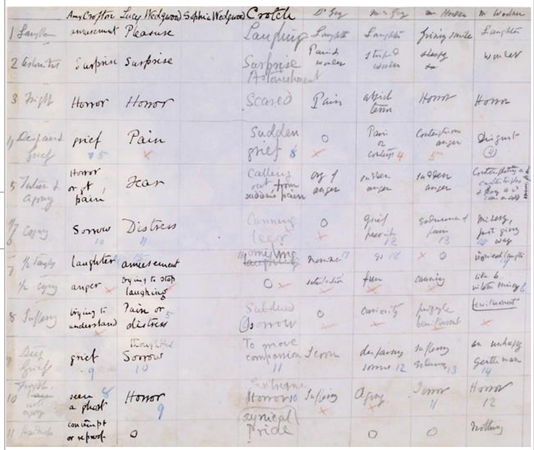
Boredom	Confusion	Fear	
			
Happiness	Sadness	Surprise	Darwin’s spreadsheet

Figure 1: Some of Duchenne’s photographs and Darwin’s spreadsheet analysing perceptions

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.

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, and particularly for work on affective computing in the past 15 years (Ekman & Friesen, 1978). 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 further 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 recognised 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 and any computation involving emotions must be designed to handle ambiguity and uncertainty (Robinson & Baltrušaitis, 2015).

Recognising emotions

Although Darwin concentrated on facial features to convey emotions in his book on emotions, he also mentions 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 six basic emotions. Recognising 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 recognising emotions that requires no human intervention and which operates in real-time on commodity hardware.

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 recognise 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. For a mean false positive rate of 5%, the overall accuracy of the system is 77%. The system also generalises well to faces not included in the training data. 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.

Evaluation separated nine conditions with an accuracy of 70%, increasing to 83% if multiple winners were considered. The approach also generalises 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. However, characteristics indicating what movement is being considered and which person is doing it must be discounted before it is possible to analyse how it is being done. 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 are beginning to 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 must 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. 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 human-likeness.

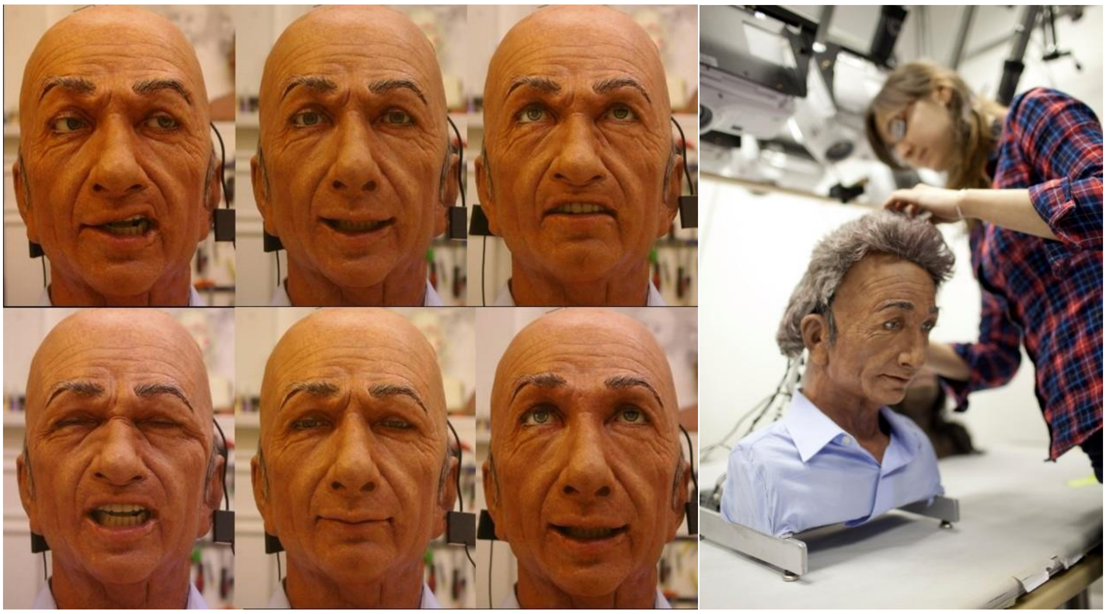

Figure 2: A humanoid robot used to examine participants' reactions to emotional expressions

The technology has been tested using a high-fidelity robotic head which simulates movement disorders such as dystonia that might be encountered by trainee doctors, and its use is currently being investigated as a possible intervention in autism therapy.

Applications

An early 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 busy 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.

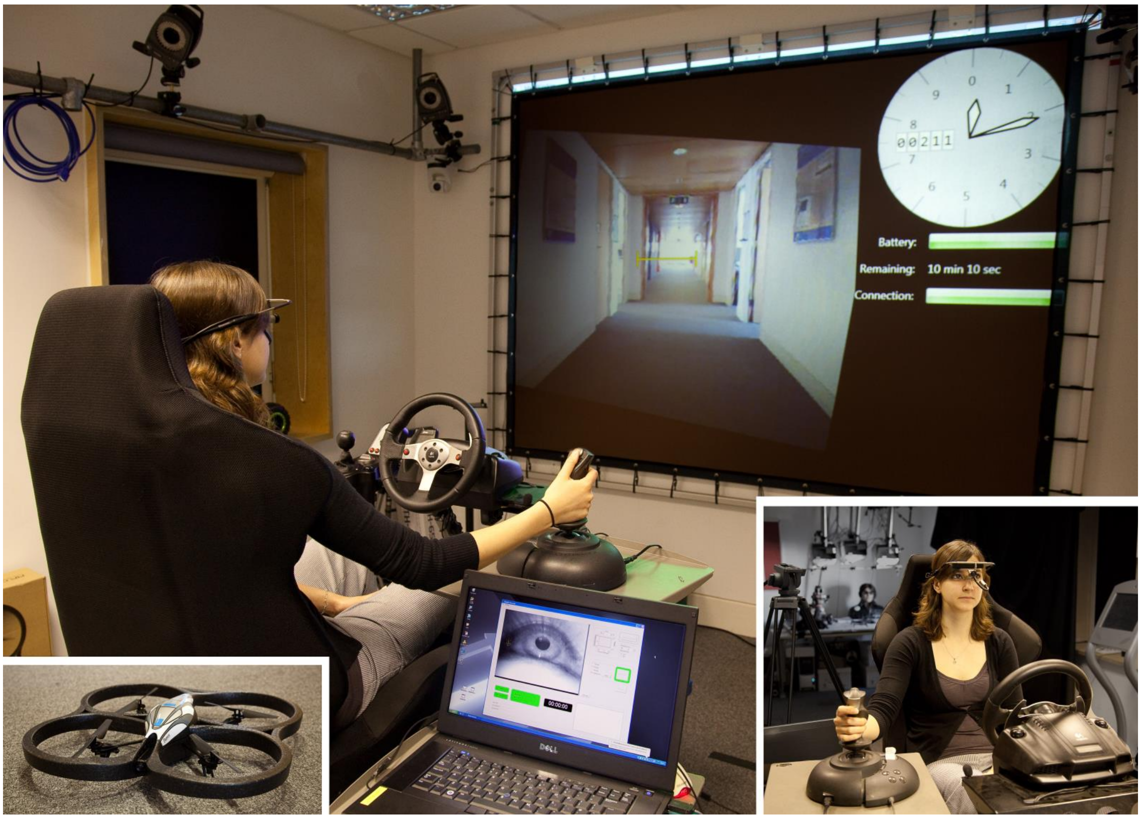

Figure 3: A simulator to investigate drivers' signs of cognitive overload

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.

More importantly, 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.

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 sorts of systems will soon be usable on portable equipment such as smartphones. This will open up many exciting applications from games to health care.

The triumph of technology

The invention of the transistor in 1947 and the production of the first integrated circuit in 1961 have been followed by 60 years of extraordinary progress. The number of transistors on a single integrated circuit 2cm square has doubled every two years, rising from 4 in 1961 to 8 billion in 2023 (Roser, et al., 2023), while the cost of the chip has remained roughly constant at a few dollars. This has made phenomenal amounts of computing power available cheaply. The volume of data available on the World-Wide Web has grown at a similar rate since its introduction in 1989. Together these advances have allowed the development of computer systems that can simulate many facets of human intelligence.

Steady progress with scientific understanding of human psychology has matched the advances in technology. Public perceptions, on the other hand, have leapt 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, the 20th century robots described by Karel Čapek (Čapek, 1921), Fritz Lang and Isaac Asimov, 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.

It may help to distinguish mechanical robots from automatic systems embedded in the world around us (Cameron, 2017). Mechanical robots include not only the homicidal cyborgs of fiction but, more importantly, industrial and domestic machines that perform boring tasks or those requiring great precision. Increasingly, they also include toys. These offer clear benefits or entertainment, but also present some challenges in changing the nature of employment.

Embedded systems, on the other hand, range from satellite navigation in cars, through internet search engines, to automated financial trading. The word 'robot' should be understood to include any system that senses its environment and can affect it autonomously (Cerf, 2013), and these hidden systems are more insidious. There are obvious dangers in the deployment of autonomous weapons by the military, but the dangers of relinquishing control of our financial systems, energy infrastructure or even cars to autonomous systems are just as worrying.

Technology moves very quickly. Recent developments and future possibilities range from robots in the work place that are already replacing blue-collar workers through to intelligent systems that might soon replace white-collar workers in professions such as medicine and law. Robotic toys are growing in popularity, both as conversational dolls for young children and as companions for the elderly. There is obvious value in these, but also a warning that we are delegating to machines the personal relationships for which we were created. Something has gone seriously wrong with our relationships when a mother post messages to her 'friends' on a smartphone while her child has a conversation with an electronic doll.

We were made in God's image to rule and for relationships, with him and with each other. Are we delegating that rule to machines? Are we sacrificing human relationships for something synthetic?

Indeed, we should think carefully about our own relationships with people and with machines. It is often convenient to anthropomorphise machines, especially when they display human aspects in their interactions or even their physical appearance. Speaking to a machine and listening to its response or interacting with a machine that happens to be wrapped up so that it that vaguely resembles a person should not fool us into thinking that it actually is a person.

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 their original sin but are a reflection of imperfections in their human creators.

This puts the nascent "robot rights" movement into perspective (Henderson, 2007). Robots are not genuinely free agents but are merely remotely controlled (possibly separated by time as well as distance) 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.

Some people are even talking about the 'singularity' when robot intelligence overtakes human intelligence (Kurzweil, 2005). This is more speculative, but still serves to keep us aware of possible dangers. However, there are immediate questions that need to be asked about autonomous weapons and social media, both of which are already having harmful effects. These dangers are just those inherent in any system created by people who are either cavalier in their motivation or lack competence. These are the general issues facing any profession but can easily be left behind with the rapid pace of change in technology.

Artificial intelligence

The term "artificial intelligence" is used very broadly: sometimes for marketing ("This vacuum cleaner has AI"), sometimes for scaremongering ("Humans face extinction as a result of AI") and sometimes just as an impressive term for computer ("The AI got my bank statement wrong). Computer scientists have used AI to refer to any problem that looked difficult and then gave it a more prosaic name when the problem was solved. It would have been helpful if Herb Simon's description of it as "complex information processing" had been adopted instead of John McCarthy's punchier "artificial intelligence".

Humanoid robots are an interesting distraction, but the real dangers lie with the computer systems embedded in the world around us that Vint Cerf identified in his notion of a robot.

Artificial intelligence is simply a machine performing a cognitive process like a person. Recent developments in AI often use machine learning. A large collection of real-world data is analysed to build a statistical model which can then be used to predict responses to unseen stimuli. This underpins recent advances such as facial expression analysis, automatic language translation, prediction of protein structures, fraud detection for credit cards, autonomous vehicles and so on.

The same technology underpins social media. Publishers like YouTube, Facebook, WhatsApp, Instagram, WeChat, TikTok and so on want to make money by selling advertising. They can sell more advertising if they can engage the user and retain their attention for longer. This involves accumulating a large body of data by recording each user's activity and using this to build a statistical model that models their interests. This can then be used to predict what else would retain their attention while they are shown more advertisements. Likewise, they can model each user's network of contacts and use that to expand their clientele. Advertising, engagement and growth.

This model of social media has had undesirable side-effects. Users are only shown material that matches their earlier interests, so their preconceptions are never challenged. More importantly the social media become a substitute for the relationships that are fundamental to humanity.

Again, fiction offers us some useful examples. *The machine stops* describes a society where people live in isolation and relationships are mediated through technology (Forster, 1909). People's attention span has shrivelled, their relationships become superficial, and they are challenged by a failure of the technology. Imagine today's world trying to survive for more than a few hours without the internet.

Brave new world describes a world in which people are reduced to passivity by a flood of information (Huxley, 1931). They no longer care about their autonomy. By contrast, *1984* describes a world in which people are deprived of information and controlled by a surveillance state (Orwell, 1949). They are no longer allowed autonomy. Modern technology presents both risks: we drown in trivia while surrendering ourselves to increasing levels of surveillance through our mobile phones.

These represent the real dangers of artificial intelligence. A large group of distinguished scientists wrote an open letter in 2023 arguing that mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war (Center for AI safety, 2023). They identify five areas of risk: the loss of human skills, bias in machine learning, election disinformation, social control and autonomous weapons. The new European AI Act aims to regulate these challenges (European Paliament, 2023) .

The image of God

Computer systems can infer a people's mental states from facial expressions, tone of voice and body posture and gesture. Computers can even display emotions through animated cartoons, physical robots and language. Do these synthetic emotions mean that they also share the image of God? The account of creation in Chapter 1 of Genesis showed God's image in His creativity, rule and relationships. But we can also see a person's image in his or her children. Children resemble their parents in their characters as well as in their physical appearance. So we see the image of God most clearly in Jesus. As Paul says in Colossians 1:15:

> *The son [Jesus] is the image of invisible God, the firstborn over all creation.*

John's gospel starts by telling us that Jesus was with God in the beginning, at the creation. And John 3:16 tells us that Jesus came to restore the relationship between mankind and God that was broken at the fall. This shouldn't come as a surprise. God loved Adam and Eve, even after they broke his rule. God declared his punishment in Genesis 3, but then it was God who made garments of skin for Adam and Eve and clothed them. The whole Bible is the story of a loving God desperate to restore the broken relationship with mankind, and mankind failing to respond.

Like father, like son, we see God's image in Jesus' death on the cross. In the words of John 3:16:

> *For God so loved the world that he gave his one and only Son, that whoever believes in him shall not perish but have eternal life.*

We see an extraordinary, sacrificial love which demands our response. You and I and all of fallen mankind can never be good enough to earn a restored relationship with God. All we can do is accept the amazing gift of God giving his one an only son to take the punishment that we deserve. If we accept that gift, we need not perish but can enjoy eternal life in a restored relationship with our creator.

There we see most clearly the difference between men and machines. God loved his creation – mankind – so much that he gave his one and only son to restore his relationship with them. I am very fond of the machines that I make, but there is no possibility that I would give my son to restore them if they were broken. In the words of the creed, it is the difference between being begotten and being made. Humans are begotten and machines are made. There is a clear difference.

The challenge for Christians is to exercise rule and enjoy relationships as God intended. Robots – in the form of autonomous systems and artificial intelligence – may have a role to play in assisting us in those, but God still expects us to rule wisely and maintain our relationships with him and with each other. We remain responsible for ensuring the motivation and competence of those building and deploying robots, and we must be careful to distinguish simulations from the real thing. Only then can we engage with changes in the world around us and bring biblical perspective to bear on them.

Our relationship with our creator God remains a fixed point as the world around us changes. And that relationship is restored through Jesus. It is what makes us human rather than automata.

Works Cited

Baron-Cohen, S., Golan, O., Wainwright, S. & Hill, J., 2002. *Mind reading: the interactive guide to emtoions,* London: Jessica Kingsley.

Cameron, N., 2017. *The robots are coming: us them and God.* s.l.:CARE.

Čapek, K., 1921. *R.U.R. (Rossum's Universal Robots).* Prague: Aventinum.

Center for AI safety, 2023. *Statement on AI Risk.* [Online]
Available at: https://www.safe.ai/work/statement-on-ai-risk

Cerf, V. G., 2013. What's a robot. *Communications of the ACM,* 56(1), p. 7.

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

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

Ekman, P., Friesen, W. & Ellsworth, P., 1972. *Emotion in the human face.* New York: Pergamon Press.

el Kaliouby, R. & Robinson, P., 2004. *Real-time inference of complex mental states from facial expressions and head gestures.* Washington DC, IEEE Computer Vision and Pattern Recognrion.

European Paliament, 2023. *EU AI Act: first regulation on artificial intelligence.* [Online]
Available at: https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence

Forster, E., 1909. The machine stops. *The Oxford and cambridge Review*.

Hawking, S. W., 2014. *Stephen Hawking warns artificial intelligence could end mankind* [Interview] (2 December 2014).

Hawking, S. W., 2018. How did it all begin?. In: *Brief answers to the big questions.* London: John Murray.

Henderson, M., 2007. Human rights for robots? We're getting carried away. *The Times*, 24 April.

Huxley, A., 1931. *Brave new world.* s.l.:Chatto & Windus.

Kurzweil, R., 2005. *The singularity is near.* s.l.:Viking.

Orwell, G., 1949. *Nineteen eighty-four.* s.l.:Secker & warburg.

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

Robinson, P. & Baltrušaitis, T., 2015. *empirical analysis of dimensional affect.* Xi'an, Affective Computing and Intelligent Interaction.

Roser, M., Ritchie, H. & Mathieu, E., 2023. *What is Moore's Law.* [Online].

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

Wajeeha, 2016. *The Russian Intelligent Robot Escapes Lab For The Second Time And Will Probably Be Scrapped.* [Online]
Available at: https://wonderfulengineering.com/the-russian-intelligent-robot-escapes-lab-for-the-second-time-and-will-probably-be-scrapped/