CHAPTER 30

Child machines

DIANE PROUDFOOT

This chapter outlines Turing's key ideas in artificial intelligence (AI) and charts his legacy in the field of robot intelligence. In 1950 Turing suggested that one approach to machine intelligence would be to provide a machine with 'the best sense organs that money can buy', and then 'teach it to understand and speak English'. After decades of struggle to create intelligent software, the current goal of many researchers in AI is indeed to build 'socially intelligent' robots—machines with vision and hearing and primitive communicative abilities. The grand aspiration of these theorists is to create what Turing called a 'child machine'—a machine that, like a human infant, can point, smile, recognize its carer's face, and learn to distinguish itself from others. In this chapter I discuss Turing's child machine and its descendants in modern cognitive and developmental robotics.

A recipe for machine intelligence

In 1950 Turing said: 'Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child's? If this were then subjected to an appropriate course of education one would obtain the adult brain'.[1] His 'guiding principle' in the attempt to build intelligent machines was to follow the development of intelligence in the human being:[2]

> If we are trying to produce an intelligent machine, and are following the human model as closely as we can, we should begin with a machine with very little capacity to carry out elaborate operations or to react in a disciplined manner to orders . . . Then by applying appropriate interference, mimicking education, we should hope to modify the machine until it could be relied on to produce definite reactions to certain commands. This would be the beginning of the process.

Turing called this simple machine a 'child machine' and said that it must learn 'initiative' as well as discipline, so that it can modify its own instructions and make its own 'choices'. When it does so, it has 'grown up'—and then 'one is obliged to regard the machine as showing intelligence'. According to Turing, this is just to follow the example of the human child: when a child learns to make discoveries independently of her teacher, the teacher does not claim the credit.[3]

Should a child machine be a disembodied 'brain' that plays chess and cracks codes, or a humanoid robot that might learn for itself by 'roam[ing] the countryside'? Turing described the disembodied and the embodied routes to building a thinking machine and suggested that researchers pursue both approaches. He said:

> It can also be maintained that it is best to provide the machine with the best sense organs that money can buy, and then teach it to understand and speak English. This process could follow the normal teaching of a child.

In his view, a child—human or machine—becomes intelligent only through education.[4]

Turing's descriptions of his child machine are frequently tongue-in-cheek. He said, for example, that the machine could not be sent to school 'without the other children making excessive fun of it' and so its education 'should be entrusted to some highly competent schoolmaster'.[5] These remarks, however, sit alongside his very serious intent—to outline a research programme for AI. For many years AI largely ignored this option, but now roboticists aim to build a machine with the cognitive capacities of human infants—a child machine. The roots of this research field in Turing's work have been neglected. In this chapter I consider how his dream has played out in developmental robotics. This also provides an insight into the challenges that face AI.

From universal machine to child machine, and back again

Turing's universal machine of 1936 can be programmed to execute any calculation that a 'human computer' can perform. But does it *learn*? For Turing, learning is the key to intelligence—in 1947 he said, 'What we want is a machine that can learn from experience'. In his view, a 'learning machine', built and educated in analogy with the education of a human child, can develop as the child does. We should:

> start from a comparatively simple machine, and, by subjecting it to a suitable range of 'experience' transform it into one which was more elaborate, and was able to deal with a far greater range of contingencies . . . As I see it, this education process would in practice be an essential to the production of a reasonably intelligent machine within a reasonably short space of time. The human analogy alone suggests this.

As it learns, the machine is to modify its own instructions—'like a pupil who had learnt much from his master, but had added much more by his own work'. Turing hoped that there would be 'a sort of snowball effect. The more things the machine has learnt the easier it ought to be for it to learn others'; the machine would probably also be 'learning to learn more efficiently'.[6]

Turing's insight was to begin with an 'unorganised' machine—a machine made up 'in a comparatively unsystematic way from some kind of standard components' and which is 'largely random' in its construction. His hypothesis, which he thought was 'very satisfactory from the point of view of evolution and genetics', was that 'the cortex of the infant is an unorganised machine, which can be organised by suitable interfering training' into a universal machine ('or something like it'). According to Turing, the structure of the child machine is analogous to the 'hereditary material' in the infant brain, changes in the machine are analogous to human genetic mutations, and the choices of the AI researcher are analogous to the influence of natural selection on humans. His goal was an unorganized machine that could be organized to become a universal machine, as a child's brain is altered by natural development and the environment.

In this process the task of the researcher is mainly to give the child machine the appropriate experiences—Turing hoped that this process would be 'more expeditious than evolution'![7]

Turing conceived of two kinds of unorganized machine. One was the first example of computing by means of neural networks—his 'A-type' and 'B-type' machines (see Chapter 29). According to Turing, the A-type machine is 'about the simplest model of a nervous system with a random arrangement of neurons', and it would not require 'any very complex system of genes to produce something like the A- or B-type'. The B-type machine is a modified A-type; a sufficiently large B-type can be trained to become a universal machine, Turing claimed.[8]

He called his other kind of unorganized machine a 'P-type' machine: this is a Turing machine with an initially incomplete program. A 'pain' stimulus is then used to cancel tentative lines of code, and a 'pleasure' stimulus to make these lines of code permanent—this procedure completes the program. In Turing's view, training a human child depends largely on 'a system of rewards and punishments, and this suggests that it ought to be possible to carry through the organising [of a machine] with only two interfering inputs, one for "pleasure" or "reward". . . and the other for "pain" or "punishment" '. The P-type was to test this hypothesis. Turing said: 'It is intended that pain stimuli occur when the machine's behaviour is wrong, pleasure stimuli when it is particularly right. With appropriate stimuli on these lines . . . wrong behaviour will tend to become rare'. He recognized, though, that education involves more than rewards and punishments, joking that 'if the teacher has no other means of communicating to the pupil . . . [b]y the time a child has learnt to repeat "Casabianca" he would probably feel very sore indeed, if the text could only be discovered by a "Twenty Questions" technique, every "NO" taking the form of a blow'. Some other 'unemotional' means of communication with the machine is required—Turing called these additional inputs to the P-type 'sense stimuli'.[9]

Turing's views on machine learning influenced others at the time, such as Anthony Oettinger, who wrote the earliest functioning AI programs to incorporate learning.[10] Oettinger's 'shopping programme' ran in 1951 on the University of Cambridge EDSAC (the Electronic Delay Storage Automatic Calculator, the world's second stored-program electronic computer). This program—which Oettinger described as a child machine—simulates the behaviour of 'a small child sent on a shopping tour'; the program learns which items are stocked in each shop in its simulated world, so that later, when sent out to find an item, it can go directly to the correct shop.[11] Also in 1951, Christopher Strachey, whose draughts-playing program (see Chapter 20) was the first to use heuristic search—part from Turing's own chess-playing program[12]—said that Turing's analogy between the process for producing a thinking machine and teaching a human child was 'absolutely fundamental'. According to Strachey, the first task is 'to get the machine to learn in the way a child learns, with the aid of a teacher'. Like Turing, he said that one of 'the most important features of thinking' is 'learning for oneself by experience, without the aid of a teacher'. Strachey believed that he had 'the glimmerings of an idea of the way in which a machine might be made to do [this]'.[13]

The computer scientist Donald Michie described himself, Turing, and Jack Good as (at Bletchley Park during the Second World War) forming 'a sort of discussion club focused around Turing's astonishing "child machine" concept'. This concept, he said, 'gripped me. I resolved to make machine intelligence my life as soon as such an enterprise became feasible'. For Michie, as for Turing, the 'hallmark of intelligence is the ability to learn' and, like 'a newborn baby', a computer's possibilities 'depend upon the education which is fed into it'.[14] In the 1960s Michie built famous early learning machines. His MENACE machine (Matchbox Educable Noughts-And-Crosses Engine) could be trained to improve its game. The FREDERICK robots (Friendly

Robot for Education, Discussion and Entertainment, the Retrieval of Information, and the Collation of Knowledge, usually known as Freddy), built in Michie's lab at the University of Edinburgh, learned to manipulate various objects, including how to put differently shaped blocks together in order to create a toy (see Ch. 25).[15] Michie later criticized AI's attempts to build human-level expert systems—programs imitating a human expert's knowledge of a specific area—on the grounds that this approach neglected Turing's child-machine concept.[16]

In 2001 Michie said that AI is part way to building a child machine, in that programmers know how to acquire and represent knowledge in a program. Now we must use these programming techniques 'so as to constitute a virtual person, with which (with whom) a user can comfortably interact'. We must build a machine that is 'a "person" with sufficient language-understanding to be educable, both by example and by precept'. The goal of AI should be, not only a human-*level* machine, but a human-*type* machine (the HAL of Kubrick's film *2001: A Space Odyssey* is the former but not the latter, in Michie's view). The teacher must have a *rapport* with the child machine. Without this rapport, the teacher is 'in effect being asked to tutor the [machine] equivalent of a brainy but autistic child'—in this interaction, Michie said, there are no 'dependable channels of communication' and the education process is unlikely to succeed.[17]

Educating a child machine

According to Turing:

> Presumably the child-brain is something like a note-book as one buys it from the stationers. Rather little mechanism [i.e. writing], and lots of blank sheets . . . Our hope is that there is so little mechanism in the child-brain that something like it can be easily programmed. The amount of work in the education we can assume, as a first approximation, to be much the same as for the human child.

Turing claimed that, 'in so far as a man is a machine he is one that is subject to very much interference [i.e. education]. In fact interference will be the rule rather than the exception. He is in frequent communication with other men, and is continually receiving visual and other stimuli which themselves constitute a form of interference'. A teacher aims to alter a child's behaviour, with the result, Turing said, that 'a large number of standard routines will have been superimposed on the original pattern' of the child's brain. The child is then in a position 'to try out new combinations of these routines, to make slight variations on them, and to apply them in new ways'. Even if a human being appears to be acting spontaneously, this behaviour is 'largely determined by the way he has been conditioned by previous interference'.[18]

A 'grown up' machine does not need so much 'interference'. Turing said:[19]

> At later stages in education the machine would recognise certain other conditions as desirable owing to their having been constantly associated in the past with pleasure, and likewise certain others as undesirable. Certain expressions of anger on the part of the schoolmaster might, for instance, be recognised as so ominous that they could never be overlooked, so that the schoolmaster would find that it became unnecessary to 'apply the cane' any more.

The educated machine has learned to generalize from past 'experience'.

Did Turing succeed in educating his child machines? He said that it should be easy to simulate unorganized machines on a digital computer; having done so, one could program 'quite definite "teaching policies" ' into the machine. 'One would then allow the whole system to run for an appreciable period, and then break in as a kind of "inspector of schools" and see what progress had been made.' However, Turing had to make do with the only programmable computers available in the 1940s—'paper machines'. These were human beings 'provided with paper, pencil, and rubber, and subject to strict discipline', carrying out a set of rules. Simulating B-types required considerable computational resources, and so was delayed until these became available—too late for Turing (see Chapter 29). He did attempt to teach a P-type, but found this 'disappointing'. He said that organizing a P-type to become a universal machine was 'probably possible', but it was 'not easy' without adding a systematic external memory, and then the supposedly unorganized machine would be more organized than an A-type. Also, Turing said, the method of training a P-type was not 'sufficiently analogous to the kind of process by which a child would really be taught' and was 'too unorthodox for the experiment to be considered really successful'. His method included letting the machine run while continuously applying the 'pain' stimulus; using this procedure, the machine 'learnt so slowly that it needed a great deal of teaching'.[20]

Turing wanted to investigate 'other types of unorganised machine, and also to try out organising methods that would be more nearly analogous' to the education of human beings.[21] But his own attempts to teach a child machine were frustrated.

The best sense organs that money can buy

Turing is often viewed as initiating a research programme to create a disembodied computer program—one that plays chess, cracks codes, and in general solves mathematical puzzles. He said that he wished 'to try and see what can be done with a "brain" which is more or less without a body, providing at most organs of sight, speech and hearing'. What will this machine do? Owing to its 'having no hands or feet, and not needing to eat, nor desiring to smoke, it will occupy its time mostly in playing games such as Chess and GO, and possibly Bridge'. How will the machine be educated? According to Turing, it would not be possible to teach the machine exactly as a teacher would a 'normal' child—for example, it 'could not be asked to go out and fill the coal scuttle'. However, the 'example of Miss Helen Keller shows that education can take place provided that communication in both directions between teacher and pupil can take place by some means or other'. To play chess, the machine's 'only organs need be "eyes" capable of distinguishing the various positions on a specially made board, and means for announcing its own moves'. Turing thought that this machine should do well at cryptography but would have difficulty in learning languages—the 'most human' of the activities a child machine might learn. Learning languages, he said, seems 'to depend rather too much on sense organs and locomotion to be feasible'.[22]

For 'locomotion', robots are required. Turing was probably the first person to recommend building robots as the route to thinking machines. In 1948 he said:

> A great positive reason for believing in the possibility of making thinking machinery is the fact that it is possible to make machinery to imitate any small part of a man . . . One way of setting about our task of building a 'thinking machine' would be to take a man as a whole and to try to replace all the parts of him by machinery. He would include television cameras, microphones, loudspeakers, wheels and 'handling servo-mechanisms' as well as some sort of 'electronic brain'. This would

of course be a tremendous undertaking. The object if produced by present techniques would be of immense size, even if the 'brain' part were stationary and controlled the body from a distance. In order that the machine should have a chance of finding things out for itself it should be allowed to roam the countryside, and the danger to the ordinary citizen would be serious.

This approach, Turing said, is 'probably the "sure" way of producing a thinking machine'. The response at the time to his plan was negative: in the view of some of his colleagues at the National Physical Laboratory, he was going to 'infest the countryside' with 'a robot which will live on twigs and scrap iron'. Turing himself said that his plan was 'too slow and impracticable' and he abandoned the idea of making a 'whole man'.[23] This idea, though, was revived in the 1990s, in Rodney Brooks's Cog project.

Living the Turing dream

Brooks, the foremost pioneer of embodied AI, has said that what drives him is the 'dream of having a thinking robot'.[24] The robot Cog had, as Turing envisaged, a 'body' (composed of a 'head', 'torso', 'arms', and 'hands'), television cameras and microphones, and an off-board 'brain'. Cog's education also proceeded in part as Turing had suggested. In Turing's view, the child machine's teacher should be someone who 'is interested in the project but who is forbidden any detailed knowledge of the inner workings of the machine'. Many people helped to train Cog, for example to reach towards objects or to play with toys; as a result, the robot's engineers were (to use Turing's words) often 'very largely ignorant of quite what is going on inside' the machine. In this respect, teaching the machine resembled teaching a human child, and was (as Turing had said) 'in clear contrast with normal procedure when using a machine to do computations'.[25]

The fields of social robotics and human–robot interaction have exploded in the late twentieth and early twenty-first centuries. One goal is to build entertainment robots—robot pets such as Pleo the dinosaur, Paro the baby seal, and the Genibo dog. Another goal is to replace expensive human labour by service robots—such as Brooks's Baxter (a user-friendly industrial robot with a torso, two arms, and a head) and Juan Fasola and Maja Matarić's Bandit (a humanoid robot that can guide elderly people through rehabilitative exercises). A robot physiotherapist has advantages other than not needing holiday or sick leave; for example, children with an autism spectrum disorder (ASD) seem to respond better to robot teachers than to the human sort. (Paro is also reported to help dementia sufferers.) In many cases, service robots must behave in human-like ways, so that human clients with no specialized training can easily explain and predict the robots' behaviour. Therapeutic robots may also be required to possess 'social intelligence'—the ability to identify and respond to their human clients' needs and desires. Without this, human–robot interaction may not work. Rapport, as Michie said, is essential.

Turing's vision of a child machine flourishes especially in developmental robotics. Here the grand goal is to build a machine that has infant-level social intelligence, acquired in infant-like ways. Typically-developing human infants follow a well-described trajectory. A child must acquire a 'theory of mind'—the concept of the distinction between herself and other people. Theory-of-mind abilities are essential to interaction with parents, siblings, and strangers. Developmental roboticists aim to construct a robot with exactly these skills—a machine that can detect faces and agents, identify what others (human or robot) are looking at and attend to the same object, recognize itself, and grasp the difference between its own beliefs and those of others. The machine is to acquire these abilities gradually, as human children do, by imitating adults and siblings. This process is also to fit with what we know of human biology and psychology; for example, researchers may design a robot arm to have a range of motion similar to the human arm, or design a robot's brain (that is, system architecture) using findings from developmental psychology and neuroscience.

Some researchers even aim to build robots that have the appearance of human infants—such as Javier Movellan's Diego-san, a hyper-realistic 'expressive' robot toddler (Fig. 30.1). The robot's creators believe that the fact that the human body is very compliant (that is, it moves as a whole, rather than as a series of independent parts) is crucial in the development of intelligence. The 'driving hypothesis' behind the Diego-san project is that humans 'are born with a deconstructed motor control system' that, as the infant matures, is reconstructed 'in a manner that leads to the development of social interaction and symbolic processes, like language'—the researchers' goal is a computational theory of this process.[26] In this respect, Diego-san is an initially unorganized child machine that is to learn as a human infant does. The robot has learned to 'reach towards objects and move them to his mouth' as a human infant does, by making high-intensity movements when the object is far away and the mother is present and then making more precise movements when the object is close.[27] This, Movellan says, is 'the very beginning of gesturing'—pointing to request an object or to draw another person's attention to the object, core theory-of-mind abilities.

Building child machines may also help us to test theories of human psychology. How do children acquire social intelligence? In this process, what are the developmental differences

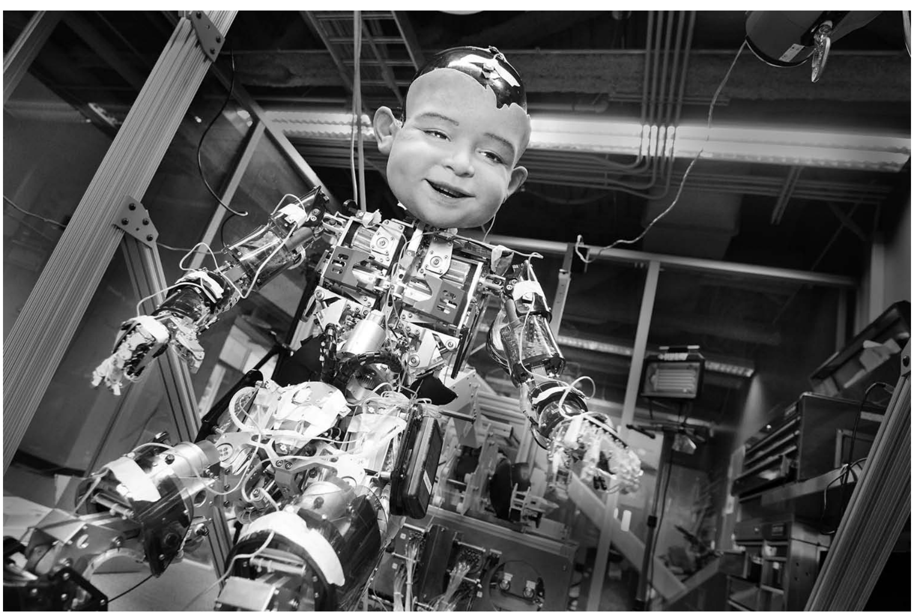

Figure 30.1 Javier Movellan's Diego-san.

Reprinted with permission of the Machine Perception Laboratory, UC San Diego; with thanks to Javier Movellan.

between the neurotypical child and the child with ASD? Researchers hope that we can investigate these and other questions by studying human–robot (or robot–robot) interaction in standardized conditions—in place of difficult and perhaps unethical studies of adult–infant (or infant–infant) interaction. This too fits with Turing's vision: he said, 'I believe that the attempt to make a thinking machine will help us greatly in finding out how we think ourselves'.[28]

Smiley faces

In modern AI 'face robots'—systems with a 'face' and 'head' and 'facial expressions'—are descendants of Turing's suggestion that we give a machine 'the best sense organs that money can buy' and teach it to 'understand and speak English' by a process following the 'normal teaching of a child'. These robots may have machinery enabling them to 'see' or 'hear', and although typically they do not 'speak English', they are intended to have earlier and more primitive communicative abilities—facial expression and bodily gesture. Faces are critical in human communication; for example, Baxter, although designed to work on a factory production line, has a 'face', in order that humans can interact with it intuitively (Fig. 30.2). Researchers engineer face robots not only to produce facial expressions but also to attend to human faces, identify human facial expressions, and respond in kind—for example, 'smiling' in response to a human's smile.[29] These abilities are the building blocks of the human infant's interaction with other humans.

According to Turing, there is 'little point in trying to make a "thinking machine" more human by dressing it up in . . . artificial flesh'.[30] Several face robots, nevertheless, are hyper-realistic representations of a human face and head. Diego-san (which has a body as well as a face) seems to have dimples. Fumio Hara and Hiroshi Kobayashi built a series of robots designed to resemble

Figure 30.2 Baxter's 'face'.

Reprinted courtesy of Rethink Robotics, Inc.; with thanks to Rodney Brooks.

the head of a young female, with quasi-realistic skin, hair, teeth—and make-up. The hyper-realistic Einstein robot in Movellan's lab, which learns to smile from initial 'body babbling' (random movements of its 'facial muscles'), has the physicist's wrinkles.[31] Turing hoped that 'no great efforts will be put into making machines with the most distinctively human, but non-intellectual characteristics such as the shape of the human body'; in his view, it is 'quite futile to make such attempts and their results would have something like the unpleasant quality of artificial flowers'.[32] This is a striking anticipation of the now much-discussed discomfort that people report when presented with a *too* human-like machine.

Human infants learn communication skills from face-to-face interaction with adults. The infant smiles spontaneously, without understanding or intending the meaning of the gesture; the adult interprets this as an expression of the baby's inner state, smiles in return, and reinforces the baby's behaviour. In this way the infant learns how to communicate—not merely to mimic the facial shape of 'smiling', but to smile. Developmental roboticists hope to teach their machines to communicate by a similar process. For example, Diego-san's 'smiling' is to replicate the smiling pattern of a 4-month-old infant. The researchers discovered that an infant times his 'smiles' so as to maximize the duration of his mother's smiling while minimizing the duration of his own smiling; when this behaviour was programmed into the robot, human observers reacted to Diego-san's 'smiling' in the same way.[33]

One face robot has received widespread public attention for its range of emotional 'facial expressions', Cynthia Breazeal's (now retired) Kismet (Fig. 30.3). Kismet's caricature child-like 'features' and 'vocalizations' (sounds resembling a baby's babbling) are designed to evoke human nurturing responses. Observers react to the robot as if it were animate—as if (when its 'head' or 'eyes' move) it is interested in what is happening around it or is looking at a specific object. Humans speak to Kismet in baby talk, apologize to the robot for disturbing it, and empathize with it by mirroring its behaviour. But is this (to use Michie's words) 'real' rapport with the machine, or merely 'illusory' rapport?[34]

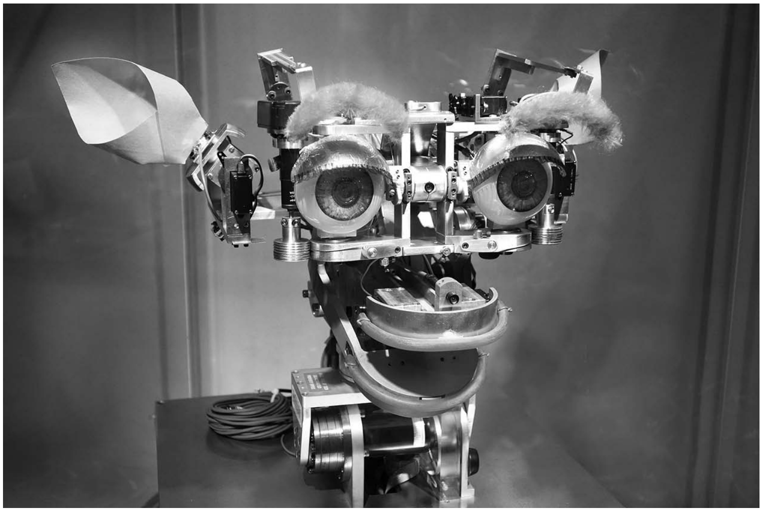

Figure 30.3 Cynthia Breazeal's Kismet.

Posted by Daderot to Wikimedia Commons, https://commons.wikimedia.org/wiki/File:Kismet,_1993-2000,_view_2_-_MIT_Museum_-_DSC03711.JPG. Creative Commons Licence.

What's in a smile?

Turing had a sophisticated approach to the concept of intelligence—an approach that is almost universally misunderstood (see Chapter 28).[35] In 1948 he said that 'the idea of "intelligence" is itself emotional rather than mathematical', and that the concept of intelligence is an 'emotional concept'. He remarked that the 'extent to which we regard something as behaving in an intelligent manner is determined as much by our own state of mind and training as by the properties of the object under consideration'.[36] In short, what makes something intelligent is not only how it behaves but also how we *respond* to it. This is not to say, though, that intelligence is solely in the eye of the observer: we may react to a machine as if it were intelligent when it is not (and so we need a test for thinking—Turing's imitation game). We may also react to a 'smile', a facial shape, as if it were a *smile*, a communicative act. Facial expressions can be understood as physical shapes or behaviours; in this sense a human can 'smile' when merely going through the motions of smiling (as when asked by a dentist to 'smile' into the mirror) or 'frown' without actually frowning (as when asked by a dermatologist to 'frown' before a botox injection). The same facial shape can be a smile or a grimace, or merely a facial tic. The chimpanzee's bared-teeth display is a very similar shape to the human smile, but it is a different action with a different meaning. So, what makes a 'smile' into a *smile*?

The additional behaviours of the 'smiler' are crucial. For example, a smile is part of a *sequence* of fluid emotional expressions; if a human being had only a 'smiling' face (or even if she had different expressions but her 'smile' were always exactly the same shape), we would not say that she was smiling. Also, a smile occurs in specific *contexts*, for example when greeting friends but not enemies; this is because we are happy to see our friends but not our enemies, and happy people too have characteristic behaviour. In addition, any smile is some *sort* of smile, for example welcoming, conspiratorial, flirtatious, or cynical; what makes a smile welcoming rather than cynical is the further actions of the 'smiler'.

Kismet does not have these additional behaviours. Its 'facial expressions' are machine-like rather than fluid, and this is very different from even the human's 'fixed' smile. When the robot 'smiles' it is not *happy*, as its creator acknowledges; just as Turing said that the words 'pain' and 'pleasure' do not presuppose 'feelings' in his P-type machines, Breazeal says that the robot's 'emotions' (that is, these components of its system architecture) are 'quite different' from emotions in humans.[37] The robot is assumed to have sorts of smiles, including a 'contented smile',[38] but in reality it lacks the behaviours that are required if a 'smile' is to be an expression of *contentment* (as when settling back into a comfortable chair and sighing peacefully after completing a task). Kismet's 'smiling' behaviour is not *smiling*—even if we react to the robot by smiling.

We all too readily anthropomorphize machines; according to Brooks, for example, human factory workers say of Baxter 'It's my buddy!'.[39] Recent films (video games, and graphic novels) that depict superhuman-level disembodied AIs and androids barely distinguishable from a human being may lead audiences to think that human-level AI is close. Misplaced anthropomorphism and make-believe are common within AI itself, and some influential but wildly optimistic researchers promise unimaginably powerful artificial minds within a few decades.[40] In reality, developmental robotics has a long way to go in order to build a child machine with infant-level social intelligence.

Growing up: discipline and initiative

Turing said, 'If the untrained infant's mind is to become an intelligent one, it must acquire both discipline and initiative . . . To convert a brain or machine into a universal machine is the extremest form of discipline . . . But discipline is certainly not enough in itself to produce intelligence. That which is required in addition we call initiative'.[41] In Turing's view a machine shows initiative when it can learn by itself, when it tries out 'new combinations' of the routines it has been given, or applies these in 'new ways'. Once it has both discipline and initiative, the child machine has grown up.

According to Turing, opponents of machine intelligence argue that a machine *cannot* think because (amongst other things) it can never 'be the subject of its own thought' or 'do something really new'. But it seems that, for Turing, the grown-up machine doesn't have these 'disabilities'. It *does* know its own thoughts: Turing said, 'At comparatively late stages of education the memory might be extended to include important parts of the configuration of the machine at each moment, or in other words it would begin to remember what its thoughts had been'. And it *can* do something new: in Turing's view, it can make 'choices' or 'decisions' by modifying its own program.[42] An AI sceptic might retort that the behaviour of the grown-up machine is merely the result of earlier programming. But, Turing pointed out, a human mathematician receives a great deal of training, which is akin to programming, and yet we don't say that the mathematician can't do anything new. Why treat the machine differently?[43]

'It is customary', Turing said, 'to offer a grain of comfort, in the form of a statement that some particularly human characteristic could never be imitated by a machine . . . I cannot offer any such comfort, for I believe that no such bounds can be set'. Smiling and other expressive behaviours are particularly human characteristics, and there is no reason to believe that a machine could never smile, or that we cannot build a socially intelligent child machine. Nevertheless, the claim that AI has already constructed 'expressive' robots is premature. And it is unlikely that Turing—who in 1952 predicted that it would be at least 100 years before a machine succeeded in his imitation game—would have said otherwise.[44]