

Genes, Affect, and Reason: Why Autonomous Robot Intelligence Will Be Nothing Like Human Intelligence

Henry Moss

Abstract: Many believe that, in addition to cognitive capacities, autonomous robots need something similar to affect. As in humans, affect, including specific emotions, would filter robot experience based on a set of goals, values, and interests. This narrows behavioral options and avoids combinatorial explosion or regress problems that challenge purely cognitive assessments in a continuously changing experiential field. Adding human-like affect to robots is not straightforward, however. Affect in organisms is an aspect of evolved biological systems, from the taxes of single-cell organisms to the instincts, drives, feelings, moods, and emotions that focus human behavior through the mediation of hormones, pheromones, neurotransmitters, the autonomic nervous system, and key brain structures. We argue that *human* intelligence is intimately linked to biological affective systems and to the unique repertoire of potential behaviors, sometimes conflicting, they facilitate. Artificial affect is affect in name only and without genes and biological bodies, autonomous robots will lack the goals, interests, and value systems associated with human intelligence. We will take advantage of their general intelligence and expertise, but robots will not enter *our* intellectual world or apply for legal status in the community.

Key words: autonomous robots, emotions, affect, artificial intelligence

“We are machines, and from that I conclude that there is no reason, in principle, that it is not possible to build a machine from silicon and steel that has both genuine emotions and consciousness.”

—Rodney Brooks, *Flesh and Machines* (2002, 180)

Cognitive scientists have begun to suggest that autonomous robots will need something similar to emotions to set goals and guide behavior. As Steven Pinker puts it, “Most artificial intelligence researchers believe that freely behaving robots . . . will have to be programmed with something like emotions merely for them to know at every moment what to do next” (Pinker 1997, 374). Without emotions, he seems to be suggesting, cognition in an autonomous machine would suffer from an avalanche of possibilities. From moment to moment, there would be a limitless number of ways to parse an experiential field and a limitless number of possible actions. Every cognitively discernible object in an experiential field would need a context-independent definition despite being embedded in a potentially infinite number of contexts defined by perspective and time, which is encapsulated as the so-called frame problem (Fodor 1983; Dreyfus 1992). Determining the best action at every moment would overwhelm a finite computational device. A freestanding machine would need a way to filter experience and quickly focus internal processing activities to generate timely and meaningful decisions and actions in a fast-changing world involving multiple and simultaneous mental states.

This problem has stimulated the development of what has been called “embodied” artificial intelligence, based primarily on modeling cognition in living systems (Calvo and Gomila 2008). The approach starts from the assumption that living systems are uniquely defined by the need to survive through maintaining self-organization and biological homeostasis in a constantly changing environment. The organism must continually appraise its environment and interact with it to promote survival through adaptation. By modeling sensorimotor feedback loops defined by such organism-environment interactions, researchers have proposed ways to ground organism behavior in terms that are concrete and not susceptible to combinatorial explosion. The maintenance of individual identity, autopoiesis, in this dynamic and transient biological milieu provides the basis for assigning intentions and therefore goals and interests to organisms (Maturana and Varela 1980). This might then be fruitfully extended to include human communication and social interaction.

Using simple organisms and sensorimotor feedback as models, embodied AI would approach robot creation from the ground up, starting with receptors and transducers linked to simple circuits or modules programmed to process information from specific stimuli. These simple circuits could then be tuned to each other in a dimensional appraisal process (Lazarus 1991). This “tuning” would generate transient, graded affective states characterizing an object or situation as, for example, more or less pleasant or unpleasant (valence) and more or less salient

(arousal level) given a programmed hierarchy of values and corresponding possible behavioral choices such as approach and withdraw.

Such a decentralized, integrated collection of specialized modules, each following a simple logic of appraisal and behavior, could approximate an organism engaged with a world. Ezequiel Di Paolo (2010), for example, suggests that organisms develop behavioral habits and preferences as adaptations that promote survival and maintain homeostasis. However, such habits would go beyond mere physical survival and define what he calls “a specific way of life” (Di Paolo 2010, 146) for the organism, providing a biologically and ecologically realistic basis for an organism’s intentions, goals, and interests. Habit formation constitutive of meaning, identity, and intentionality could ground an autonomous machine (Di Paolo 2010, 147–55).

Embodied artificial intelligence has generated a wealth of research into the dynamics of living systems and robotics labs have been designing and building machines based on these dynamics. Simple invertebrate robots have already been built based upon the dynamics of sensorimotor feedback loops and some labs have moved on to designing artificial animals including mammals (Calvo and Gomila 2008, 79–158; Hoffman and Pfeifer 2012).

Embodied Social Robots

The work of Di Paolo and others has provided important insights into the building of autonomous agents using natural principles, but researchers looking beyond simple habit toward higher forms of cognition have found challenges trying to fit these capacities into situational and concrete behavioral contexts. Wilson suggests that studying how body behavior is directly and indirectly involved in voluntary control, analogical reasoning, and imitation may point the way to modeling reasoning and abstract thought (Wilson 2008).

Part of the challenge, however, is the seeming necessity to use functional approaches and symbol processing where decision-making involves intelligent choice. Wendell Wallach and Colin Allen see this as necessary for moral reasoning in autonomous robots (Wallach and Allen 2009). Cynthia Breazeal and Rodney Brooks, among others, have been designing interactive software for social robots aimed at giving them human-like characteristics including functional emotions and a hierarchy of programmed dispositions and criteria for choosing among behavioral options (Breazeal and Brooks 2005, 274; see also Picard 1997). These software modules “play an important role in signaling salience, to guide attention toward what is important and away from distractions, thereby helping to

effectively prioritize concerns” (Breazeal and Brooks 2005, 274). The cognitive system “is responsible for interpreting and making sense of the world,” whereas the emotion system “is responsible for evaluating and judging events to assess their overall value with respect to the creature (e.g., positive or negative, desirable or undesirable, etc.)” (Breazeal and Brooks 2005, 274; see also Rolls 1999). These evaluations would provide the needed situational or embodied dimension.

In response to critics who say that such emotions are only superficial simulations, Breazeal and Brooks say they are not claiming that robot emotions would be indistinguishable from their correlates in animals and humans, but that they are not therefore trivial or “fake”: “They serve a pragmatic purpose for the robot that mirrors their natural analogs in living creatures,” making them “function better in a complex, unpredictable environment” (Breazeal and Brooks 2005, 276). To emphasize their close ties to natural emotions, however, they call their proposed artificial emotions by their human names including artificial disgust, fear, anger, happiness, and sorrow.

Implicit in this thinking is that since humans are autonomous machines with built-in affective systems, including emotions, to allow a robot to think and reason like humans we will need to incorporate *human-like* emotions into robots. Breazeal and Brooks explicitly refer to their approach as a “functional perspective.” The functional role of emotions, they say, makes human-like emotions both conceptually and technologically possible. They consider the human emotion of disgust, for example, to be a “reaction to unwanted intimacy with a repellent entity” and their Kismet robot would then show artificial disgust by “signaling rejection of an unwanted stimulus” (Breazeal and Brooks 2005, 294–95).

Such artificial emotions are shallow, however, and only faintly resemble real emotions. Part of the problem stems from the limited biological perspective of many embodied approaches to cognition. While physical survival and maintenance of homeostasis are obviously important for organisms (Di Paolo calls survival the “mother value of all values” [Di Paolo 2010, 146]) and robots may “worry” about a low battery, most human social emotions are, in fact, associated with the biological drivers of reproduction, including those related to courtship, sex, status, mating, and childrearing. For humans, such emotions heavily impact social life and culture and reflect an evolutionary logic (Kelley 2005). That is, there must be coherence between the perceived goals of an individual organism concerning courtship, sex, childrearing, and social status, for example, and the underlying biological adaptations that promote both survival and successful reproduction.

Can freestanding robots have emotions that mirror natural affect despite the linkage of such affect to underlying biological dynamics? It is important to note that biological drives do not always align with the interests of a particular organism at a particular time. Once our own bodies have exhausted their reproductive potential, for example, they are pre-programmed to decay, as when calcium metabolism needed for bone growth in youth causes ossification and brittleness in the frail aged, or when the aging body fails to adequately compensate for programmed cell death. Natural selection errs on the side of strengthening the organism during its period of reproductive potential even if it is left vulnerable thereafter. In some lower organisms, male death may follow immediately upon copulation, as when some male spiders and praying mantises offer themselves up as food for females as an unconscious, evolved method for sustaining offspring and the genes that inhabit them. Female salmon exhaust themselves to the point of death to the same end.

While human males do not commit suicide after copulation, such biologically-based behavior gives pause. A powerful underlying biological driver of organism behavior relates to achieving reproductive success. Even physical survival and maintenance of homeostasis can be affected by these dynamics as evidenced by suicidal insects and fish as well as the immense social cost of maintaining post-reproductive human organisms. Human affect evolved in this context and thus it should not be surprising to see continuity in the physiological mechanisms at work, from the limbic brain system to the dynamics of ancient hormones, pheromones, and neurotransmitters. The difference lies mostly in the intensely social nature of human affective-cognitive dynamics.

The artificial intelligence community, classical and embodied, however, operates largely on the assumption that, regardless of the evolutionary processes that brought human affective and cognitive systems into existence, by focusing on the functions and purposes of affect, emotions, and the behavioral outcomes they make possible, we can provide artificial versions that promote effective human-robot social interaction and associated intelligent robot decision-making (Breazeal and Brooks 2005; Sloman et al. 2005). Embodied approaches to AI, while rejecting classical AI and purely cognitive appraisals, are still bound by behavioral outcomes that fit the interactive models. Generating “appropriate” behaviors in a given context is the ultimate goal, implying that actual biology can, ultimately, be side-stepped.

This approach fails to recognize the deep ties that human affect has to human biology, however, and how affective systems shape cognition, reasoning, and behavior uniquely, and in ways that change significantly across the developmental

phases of a human lifetime. While there is value in defining a relationship between a specific, functionally-defined, emotion and a behavior or set of behaviors, we must be aware that the functional definition will necessarily miss important aspects of the emotion, including its social role. This, in turn, will preclude functional equivalence. Artificial affect and emotions are possible, but the goals, interests, and values they define and promote in autonomous robots are not like those of biological humans, despite the confusing terminology employed by robot theorists. The intelligence and adaptive behavior such robots display, much of which will no doubt benefit humans, will be distinctly unlike that of the humans that build and interact with them. They will obey our commands and provide expertise, but will not enter *our* intellectual world, pass any meaningful Turing test, or apply for legal status in the community.

By focusing on how affect and associated cognition, reasoning, and behavior relate to the origins and foundation of human biology, we can identify major categories of emotions where there will be no true machine equivalent and only superficial analogs. These emotions, in turn, are intimately involved in the way humans think.

The Intelligence of “Empathic” Robots

Those seriously pursuing the development of humanoid robots recognize that such robots will not be mothers, fathers, spouses, children, boyfriends, and girlfriends. They will not become sexually aroused, romantically attracted to another being, or engage in courtship and sex. They will not meaningfully substitute for a mother during infancy and childhood. They will not experience the joys and travails of being a teenager. They will not enjoy hip-hop music, fancy a certain hat, or shed a tear at a Puccini opera. They will not be expected to engage in conversation or pass a realistic Turing test involving these kinds of matters.

Fundamental human goals, interests, and values cannot be considered in the abstract and human intelligence involves solving problems and making decisions relating to goals, interests, and values that are particular to the species *Homo sapiens*. This intelligence is not rational, per se. There are no right and wrong answers, only degrees of success in reaching goals. Indeed, some goals may even conflict with others such as when a natural desire to be successful in sports leads to a decision to negate a natural inclination to have children.

The goals and values associated with courtship, romance, sex, mating, and childrearing are universal for the species and have defined much of social life since the period of early humans. Biological wiring helps to define, for example,

the desires, dispositions, and motivations of mothers and children through a complex and dynamic physiology of hormones, pheromones, and neurotransmitters. These circuits and modules guided mothers and children in the early period of human evolution and they are still alive and well today. They mediate reasoning and decision-making about family, relationships, and social status.

Robots will not have the biological drivers of humans, nor the associated perceived goals of, for example, courtship, sex, mating, and childrearing. They may be aroused, but they will not be sexually aroused. They may form attachments, but not mother-child attachments. This is why robot intelligence will be unlike human intelligence.

Yet, autonomous robot developers claim that machine affect and emotions need not relate directly to flesh and blood biological systems to be “human-like” in important ways. Machines may not get sexually aroused or love their grandmother, but there may still be functional aspects of affect and emotions that can provide the basis for intelligent machine behavior that can be called humanoid. Artificial emotions in social robots are of particular interest.

Social robots interact with humans and can “read” the cognitive element of the emotions of the engaged human as part of its evaluative apparatus. Such so-called “empathic” robots could then advance the goals of both humans and themselves (Breazeal 2002). Perhaps they will be aides, companions, teachers, or nursemaids who gain satisfaction from the well-being of the client. The robot could find “pleasure” in advancing a meaningful goal of humans such as monitoring, comforting, or teaching a child. The robot could engage in meaningful interaction with both parent and child, effecting a social bond. An embodied robot could also get “frustrated,” “embarrassed,” or “disappointed” at times when it does not receive success signals, but could learn from its experiences and adapt behavior to improve its skills, solve problems, and meet changing goals in a dynamic environment. Parents and children can, in turn, react empathically to the robot’s emotional state, creating a two-way social relationship.

Does this bring us closer to human intelligence? Despite appearances, such robots and such interactions would not be addressing the broader interests, values, and goals of humans. Instead, they would be dealing with select, narrow sub-goals chosen by the programmer. These sub-goals would be defined outside of real life circumstances and would ultimately be restricted by the cognitive formalism of the program, usually an appraisal framework. A sub-goal might be to ensure the safety of a child while at play. The robot would assess the environment for potential dangers and continually evaluate the risk level of the child’s behavior. Since

running is a common source of falls, the robot could be programmed to gently call out, “Stop running, you might get hurt!” This might remind the child of a prior warning from a parent or teacher. The child could then look at the robot as *loco in parentis* and initiate a meaningful dialogue about dangers and limits, just as she might with a human teacher.

But there would be no true “empathy” with the child. The child enjoys running, smiling as she goes. It is an important aspect of play and helps build strong, healthy children who will benefit as they transition into the teen years. A mother intuitively understands this. She empathizes and tries to balance safety with fun and healthy play. While the robot may seem to empathize with a mother’s worries, neither the mother nor the child are likely to view the robot as engaged in the same kind of emotional evaluation. The robot evaluation is strictly cognitive. If a programmer elected to punch in the full range of experiences and options posed by the seemingly simple relation between childhood play and risk, she would quickly experience the challenge of cognitive overload. Combinatorial explosion would raise its ugly head as a horde of children enthusiastically attacked a local playground. Navigating such a complex situation, based upon multiple and conflicting affective dynamics, would overwhelm an autonomous robot, but would be routine for a teacher or parent.

The same kind of problem will arise with robots that teach and serve as nursemaids or companions. They can use heuristics to assess and meet needs, diagnose and solve technical problems, and advance client well-being, but this will not involve emotional empathy with humans. If a goal in having a child attend school is to have her learn mathematics, a robot teacher can meet this need and do so without emotional involvement. If a further goal in having a child attend school is to enhance her social status and to ultimately ensure access to mates with resources, something implicitly appreciated by parents and most human teachers and associated with a biological imperative, the robot teacher will indeed need to understand and internally appreciate the affective basis of this underlying goal in order to effectively empathize with a child who may very well hate the unnatural regimentation of school including the study of mathematics. This, however, is beyond the scope of machines.

Being a friend, teaching children, and composing art involve powerful instincts, feelings, and emotions and are part of the dynamic of affective and cognitive systems of biological humans and their families, communities, and culture. Will even the most brilliant autonomous robot effectively calibrate its emotional relationship to an elder suffering from progressive dementia or approaching death

from a terminal disease? Will it be able to effectively interact with the family members of the elder who may be experiencing conflicting emotions about such circumstances? Will a freestanding robot be able to see the social world from a child's perspective, a unique perspective that changes in many ways during the course of development? A poet writes for an explicit or implicit human reader whose emotions he wishes to engage. Will a robot be able to emotionally engage such a reader or follow the thinking of a poet?

Empathy Is a Two-Way Street

Robots cannot truly empathize with humans because they do not have the biological machinery that specifies particular emotions. If it is assumed that human emotions are built on a basic logic of pleasure-pain/attraction-repulsion and involve common neurological circuits in, say, the hypothalamus or amygdala, one would still be at pains to specify the emotion as expressed. For example, no cognitive representation provides the basis for empathizing with human sexual arousal. The same will be true for disgust, anger, fear, jealousy, grief, or any other biologically-based emotion since they arose to address natural human contexts. Anger and jealousy were originally reactions to an actual or potential loss of social status, inhibiting the potential for attracting a desired mate. Disgust, defined by somatic responses like nausea or gagging, evolved to prevent humans from ingesting potentially harmful toxins, especially animal excrement and detritus. Evolutionary psychologists have done much to trace the specificity of these and other emotions to selection pressures in primate and hominid evolution (Pinker 1997; Tooby and Cosmides 2000). Neurobiologists are finding distinct neural systems that specify emotional states, including the effects of select neurotransmitters and hormones.

Humans can empathize with other humans, because we have a common set of evolved biological affective systems. We can even empathize with higher mammals that we suspect can feel pain. While the neurological mechanisms for empathy have not been fully pinned down, some think that mirror neurons in the cerebral cortex are involved as triggers for affective reactions in social species (Gallese 2009; Ramachandran 2011).

Michael Arbib recognizes the problem facing those seeking to design empathic robots (Arbib 2005, 373; see also Arbib and Fellous 2004). He notes that to the extent that a human biological emotion is built on an evaluation such as reward and punishment, a robot would need a basis in the same evaluation in order to express empathy. A computing device aiming to teach, for example, "may use a model of the student's emotions, yet may not be itself subject to . . . reward and

punishment” (Arbib 2005, 373). The teacher could not empathize with the student because the simulation is “purely of the ‘other,’ not a reflection of the other back onto the self” (Arbib 2005, 373).

In some circumstances, a robot might “read” a mental representation of a human by interpreting a facial expression, action, or general behavior. If equipped with parallel representations, the robot could then act appropriately in light of the assumed representation. Such an appraisal-based approach could certainly then lead a robot to engage in behavior useful to a human, but, as Arbib points out, “it might be a matter of terminology as to whether or not one would wish to speak of such a robot as having emotions” (Arbib 2005, 374). He further points out that robots will necessarily face situations in which there is nothing to “read” in a human’s face or behavior and asks: “Can one, then, ascribe emotions to a robot . . . for which empathy is impossible?” (Arbib 2005, 377).

Arbib suggests that “perhaps a more abstract view of emotion is required if we are to speak of robot emotions” (Arbib 2005, 377). He goes on to propose that we identify an “ecological niche” of robots wherein robot “emotions” are defined independently of human biological emotions. They would enable robots “to effectively achieve their own goals” (Arbib 2005, 377).

The Limits of Functionalism: Artificial Disgust?

In agreement with Arbib, this paper sees autonomous robots as occupying their own niche and meeting their own goals even if these goals are to serve humans. He suggests that we might use the term “motivational systems” rather than “robot emotions” to make clear that “not all emotions need be like human emotions” (Arbib and Fellous 2004, 554). Cognitive appraisals carried out within parameters defined by sensors and information processing can be enormously useful, creating mobile and flexible expert systems. They can certainly help organize behavioral responses. But if such appraisals are simply cognitive, we find ourselves once again confronted with combinatorial explosion and the frame problem, creating sharp limits on what a robot can do in real-time social circumstances. Purely cognitive appraisals would need to operate at a very simplistic level of human social intelligence.

Are there *any* biologically-based human emotions that allow for a functional approach in autonomous robots? There are limits to functionalist and behaviorist approaches, to a greater or lesser degree, in all human emotions, not just those directly associated with courtship, sex, mating, and child rearing. Competition for individual or communal resources, for example, involves risks as evidenced by

high levels of violence associated with these activities in early societies. This history is reflected in our innate inclination toward anger, jealousy, embarrassment, and fear in social situations where we bump up against others and are forced to defend our turf.

Even emotions that appear to be strictly aimed at the survival of the individual have a broader social dimension, including affective systems selected because they promoted survival through avoidance of predators and toxins. As noted above, in their plan for developing the humanoid Kismet robot, Brooks and Breazeal selected a cluster of artificial emotions, which due to their functional value might serve an empathic robot. This led them to distill from specific human emotions what they consider a functional core. As previously mentioned, they consider the human disgust emotion to be a “withdrawal reaction to a repellent entity” and Kismet would show artificial disgust by “signaling rejection of an unwanted stimulus” (Breazeal and Brooks 2005, 294–95). Biology is neatly sidestepped, but are we really dealing with the disgust emotion here as well as the human cognition and intelligence associated with the circumstances in which it arises?

Disgust, an instinct with a long evolutionary history, relates to genetic variations that protect some animals, including humans, from environmental toxins (Curtis 2013). For the instinct to work, the toxic threat would need to trigger a physiological reaction that would prevent ingestion or force regurgitation, typically a sensation of nausea or a gagging reflex. Odors, tastes, textures, and even appearances trigger the reaction, which might prevent a human from ingesting animal or human feces, vomit, decaying flesh, or a toxic plant. Disgust originated in relation to food and the digestive system. An encounter with a disgusting substance may automatically cause one to shrink back (withdraw) and contort one’s face to enhance the response and make it easier to convey to others, particularly the young.

The disgust emotion is universal and powerful. A permanent food aversion can be caused by one bad experience. It should not be surprising, then, to see humans use the term, and mimic the expression, to communicate a strong opinion in non-biological contexts, serving as emphasis in showing disdain for someone’s behavior, for example. But such uses do not involve nausea or other somatic reactions associated with food and digestion.

A disgust emotion for Kismet, with no digestive system or food aversions, would be an emotion in name only and any programmed assessment and reaction to an “unwanted stimulus” would be only a simple cognitive circuit, with no true affective component. If Kismet was to be “disgusted” by, say, animal waste, it

would tend to withdraw. However, if Kismet's programmed goal is to pick up and remove the detritus, it would have to be "attracted" to the detritus and motivated to approach it, pick it up and dispose of it. This attraction is referred to as "interest" in the Kismet taxonomy of emotions (Breazeal and Brooks 2005, 296) and would, presumably, override the withdrawal tendency much like a parent who signals disgust and yet changes a child's diaper. The robot would then consider a successful outcome to be most pleasant, a "joy" emotion in the same taxonomy, signaling, presumably, effective social interaction. A human would be pleased by the outcome as well.

In this case, it is the human who experiences the disgust emotion and tends to withdraw. If the ultimate goal of the robot is to enhance the well-being of humans, however, it might simply be programmed to identify detritus and dispose of it, much like a vacuum cleaner robot is programmed to find and eliminate dust balls. It is the identification skill and disposal ability that creates an interest in humans for waste-disposing robots. If the robot associates a human facial expression with the detritus, this would be a cognitive enhancement of its program and not emotional empathy. Why would the robot feel the urge to withdraw from a stimulus when the point is to approach?

The robot can be an expert in biological waste disposal, but there would be no natural analog or functional equivalence to a human disgust emotion. The choice of "disgust" to describe such a cognitive robot state would be arbitrary and misleading.

Even in this narrow domain, it is clear that the instinctive, biological disgust emotion integrates with human cognition and thought in special ways. For example, all cultures have a cleanliness value which informs social relations and communication. Although details may differ, most cultures put animal waste at the top of the list of things that should be isolated and either disposed of or recycled. It has been shown that even young children have innate responses that align them with this value at a basic level. We see this when a child refuses to drink with an insect in the glass, or gags on slimy, but otherwise healthy, fruits and vegetables.

Parents find it easy to teach cleanliness to children by associating dirt, or dirty things, with these innate aversions and it is not surprising that early food aversions have a strong tendency to continue into adulthood. This blend of affect, reinforcing associations, and cognition has ramifications for how we think about and carry out basic survival functions such as foraging, food storage, cooking, eating, and toileting. These serve to shape human economics and culture in many ways.

Without real biological emotions, robots would lack the necessary somatic filters and would have to approach the survival value of cleanliness at a purely cognitive level following programmed waste disposal rules. The rules would have to be fairly simple and geared to human needs. The machine would be hard-pressed, for example, to evaluate potentially complicated situations that parents evaluate quite naturally, such as deciding if a piece of a food that drops to the floor can be eaten. As a purely “cognitive” task, the specter of cognitive explosion lurks for an autonomous robot. One man’s garbage could be another man’s best meal.

Conclusion

We love our pets and empathize when they are hurt or hungry. They are part of human culture and significant resources are deployed to maintain them. Many would consider their pet to be a friend and may cry when they die. Laws are passed to prevent the abuse of pets and other animals. But pets do not think like humans and do not have human intelligence. They cannot engage in social communication that would secure equal status in a human community. The same will be true for autonomous robots. Detached from biological needs, robots will be missing true affect and the kind of decision-making and behavior that might form the basis for achieving standing in our community or answering some very simple questions in a Turing test, questions that a mother or child could answer easily.

Some argue that we are only at the dawn of the age of machine intelligence and that computer engineering will eventually add the missing dimensions by creating artificial neurons and brains or by adding the needed somatic components including the functional equivalents of neurotransmitters, pheromones, and hormones. After all, they suggest, the effects of these bio-chemicals are ultimately expressed through brain states. This could take up to a few centuries, perhaps, but the trend line has been established and the technological development needed, though difficult, is tractable. We already know what purely symbolic cognitive systems can do with respect to chess, checkers, and *Jeopardy!* Humans are machines, so it will just be a matter of time (Chalmers 2010; Kurzweil 2005).

Humans are physiochemical machines. However, if such machines are beyond biology, as Ray Kurzweil suggests, they will not engage in courtship, mating, and childrearing. They will not have *human* intelligence.

References

- Arbib, Michael. 2005. "Beware the Passionate Robot." In *Who Needs Emotions?*, ed. Jean-Marc Fellous and Michael Arbib, 333–83. New York: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780195166194.003.0012>
- Arbib, Michael, and Jean-Marc Fellous. 2004. "Emotions: From Brain to Robot." *Trends in Cognitive Sciences* 8(12): 554–61. <http://dx.doi.org/10.1016/j.tics.2004.10.004>
- Breazeal, Cynthia. 2002. *Designing Sociable Robots*. Cambridge, MA: MIT Press.
- Breazeal, Cynthia, and Rodney Brooks. 2005. "Robot Emotion: A Functional Perspective." In *Who Needs Emotions?*, ed. Jean-Marc Fellous and Michael Arbib, 271–310. New York: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780195166194.003.0010>
- Brooks, Rodney. 2002. *Flesh and Machines: How Robots Will Change Us*. New York: Random House.
- Calvo, Paco, and Toni Gomila, eds. 2008. *Handbook of Cognitive Science: An Embodied Approach*. Oxford: Elsevier.
- Chalmers, David. 2010. "The Singularity: A Philosophical Analysis." *Journal of Consciousness Studies* 17: 7–65.
- Curtis, Valerie. 2013. *Don't Look, Don't Touch: The Science behind Revulsion*. Oxford: Oxford University Press. <http://dx.doi.org/10.7208/chicago/9780226089102.001.0001>
- Di Paolo, Ezequiel. 2010. "Robotics Inspired in the Organism." *Intellectica* 1–2 (53–54): 129–62.
- Dreyfus, Hubert. 1992. *What Computers Still Can't Do*. Cambridge, MA: MIT Press.
- Fodor, Jerry. 1983. *The Modularity of Mind: An Essay on Faculty Psychology*. Cambridge, MA: MIT Press.
- Gallese, Vittorio. 2009. "Mirror Neurons, Embodied Simulation, and the Neural Basis of Social Identification." *Psychoanalytic Dialogues* 19: 519–36. <http://dx.doi.org/10.1080/10481880903231910>
- Hoffman, Matej, and Rolf Pfeifer. 2011. "The Implications of Embodiment for Behavior and Cognition: Animal and Robotic Case Studies." In *The Implications of Embodiment: Cognition and Communication*, ed. Wolfgang Tschacher and Claudia Bergomi, 31–58. Exeter: Imprint Academic.
- Kelley, Ann E. 2005. "Neurochemical Networks Encoding Emotion and Motivation: An Evolutionary Perspective." In *Who Needs Emotions?*, ed. Jean-Marc Fellous and Michael Arbib, 29–77. New York: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780195166194.003.0003>
- Kurzweil, Ray. 2005. *The Singularity Is Near: When Humans Transcend Biology*. New York: Penguin.

- Lazarus, Richard. 1991. *Emotion and Adaptation*. New York: Oxford University Press.
- Maturana, Humberto, and Francisco Varela. 1980. *Autopoiesis and Cognition: The Realization of the Living*. Dordrecht: D. Reidel Publishing.
<http://dx.doi.org/10.1007/978-94-009-8947-4>
- Picard, Rosalind. 1997. *Affective Computing*. Cambridge, MA: MIT Press.
<http://dx.doi.org/10.1037/e526112012-054>
- Pinker, Steven. 1997. *How the Mind Works*. New York: Norton.
- Ramachandran, Vilayanur S. 2011. *The Tell-Tale Brain: A Neuroscientist's Quest for What Makes Us Human*. New York: Norton.
- Rolls, Edmund T. 1999. *The Brain and Emotion*. Oxford: Oxford University Press.
- Sloman, Aaron, Ron Chrisley, and Matthias Scheutz. 2005. "The Architectural Basis of Affective States and Processes." In *Who Needs Emotions?*, ed. Jean-Marc Fellous and Michael Arbib, 203–44. New York: Oxford University Press.
<http://dx.doi.org/10.1093/acprof:oso/9780195166194.003.0008>
- Tooby, John, and Leda Cosmides. 2000. "Evolutionary Psychology and the Emotions." In *Handbook of Emotions*, 2nd edition, ed. Michael Lewis and Jeannette M. Haviland-Jones, 114–37. New York: Guilford.
- Wallach, Wendell, and Colin Allen. 2009. *Moral Machines: Teaching Robots Right from Wrong*. New York: Oxford University Press.
<http://dx.doi.org/10.1093/acprof:oso/9780195374049.001.0001>
- Wilson, Margaret. 2008. "How Did We Get from There to Here? An Evolutionary Perspective on Embodied Cognition." In *Handbook of Cognitive Science: An Embodied Approach*, ed. Paco Calvo and Toni Gomila, 375–93. Oxford: Elsevier.
<http://dx.doi.org/10.1016/B978-0-08-046616-3.00019-0>