

What baboons, babies and Tetris players tell us about interaction: a biosocial view of norm-based social learning

STEPHEN J. COWLEY[†] and KARL F. MACDORMAN^{*‡}

[†]School of Psychology, University of Hertfordshire, Hatfield, Hertfordshire, AL 10 9AB, UK, and
School of Psychology, University of KwaZulu-Natal, 4001 Durban, South Africa

[‡]School of Informatics, Indiana University, Indiana 46202, Indianapolis, USA

(Received 1 February 2006; in final form 20 June 2006)

Could androids use movements to build relationships? For people, relationships are created with the help of behaviour-shaping norms, which infants begin to discover and manipulate by the third month. To build relationships, machines can also learn to exploit human reactions in real-time decision-making. In the video game Tetris, for example, affect co-opts computer-generated patterns to simplify cognitive tasks: norms mediate what Kirsh and Maglio (*Cognitive Sci.*, 18, pp. 513–549, 1994) term *epistemic actions*, which allow implicit knowledge to shape key pressing in ways that, given past games, are likely to be informative and valuable. Expert players act to change their cognitive states by allowing the game's higher-level states to constrain their lower-level actions. Since this process enables the development of expertise, we might expect it to be widespread; but it seems marginal in hamadryas baboons, although they use affect and complex norms. In humans, by contrast, infants use adults as cognitive resources in developing their epistemic abilities. This has engineering implications for android designers. Since androids can elicit epistemic actions, engineers need to develop an affect-sensitive interface. If successful at this, even rudimentary co-action may prompt people to report experiencing androids as both making choices and violating expectations.

Keywords: Cognitive development; Epistemic action; Distributed cognition; Human–robot interaction; Intersubjectivity

1. Why start with norms?

Early proponents of artificial intelligence (AI) set out to simulate disembodied, abstract reasoning, instead of flexible, adaptive behaviour. Historically, this tradition arose because computing machines are able to process symbolic representations in ways that are systematic, generative and inferential (Fodor and Pylyshyn 1988). This kind of processing is possible because the syntax of a representation encodes its role in inference. The drawback of these *symbol systems* is that their workings depend solely on internal constraints, namely, their syntax; but internal constraints are not enough because symbol systems must also be sensitive to the ecological relationship between the agent and the world as it changes over time. In his 'Chinese room' argument, Searle (1980) argued that symbol systems are incapable of understanding: someone

could use a rulebook to produce correct Chinese output from correct Chinese input without understanding a word of Chinese.

Harnad (1990a) identified a problem with symbol systems that does not rely on arguments based on first-person phenomenology: while the symbols in these systems may be defined in terms of other symbols, they are never actually connected to the objects, events and states of affairs they are supposed to represent in the world (also see Bickhard and Terveen 1995). People often fail to notice this linking problem when interacting with a symbol system through word-forms, because they themselves ground symbol labels effortlessly, based on past experience of interacting in the world (Harnad 1990b). Thus, even Eliza and other simple text matching programs could produce the illusion of intelligence during short exchanges (Weizenbaum 1966). New AI has avoided these kinds of tricks in its quest for intelligence by focusing on embodiment (Lindblom and Ziemke 2006); but even for New AI, conclusions about what has been achieved depend on the observer's attributions. A clever designer can create a machine that is capable of such apparently purposeful behaviour as collecting aluminium cans without any internal representation of a goal (Brooks 1991). The task of locating, approaching and grasping the can is broken down into separate sensor-actuator units that function without central control. The current emphasis on situated, embodied action has promoted the development of androids that closely mimic human appearance and movements (MacDorman and Ishiguro 2006a, b, Ishiguro 2006). A skilfully crafted humanlike android can create a stronger impression of human presence than any text-based interaction. However, according to sceptics, bodily illusions of intelligence have replaced those of symbols, and indeed a problem common to some of the most impressive systems is that they focus on the exhibition of intelligent behaviour, not its cognitive development (Ziemke and Lindblom 2006, Lindblom and Ziemke 2006, M. B. Cowley 2006). At such a time, *android science* needs machines capable of more than mindless imitation.

Could androids be built that are 'capable of sustaining natural relationships with people'? (MacDorman and Ishiguro 2004.) In seeking an answer, social intelligence becomes of paramount interest to engineers. In biology, hierarchical systems self-organize as a consequence of managing body-world activity. If androids are to behave mindfully, they too will use the world beyond their bodies to soft-assemble control over action and perception. Traditional conceptions of human information processing in cognitive science have supported the sense-model-plan-act framework in robotics, in which task performance is divided into distinct stages that are performed in series. Most physical actions are conceived as being *pragmatic*, that is, they bring the world closer to a desired state of affairs. In robotics, this architecture has been under attack since at least the late 1980s for numerous reasons (Brooks 1991, MacDorman 1999). There are too many delays between sensing and action; creating a purely symbolic model of the world that is adequately rich creates a huge computational bottleneck in the information flow; and it is unclear how the system can avoid reasoning about relatively stable aspects of the world (*i.e.* the *frame problem*).[†]

Kirsh and Maglio (1994) demonstrated that human brains also do not abide by the sense-model-plan-act framework. In playing the video game Tetris,[‡] for example, actions are often *epistemic*. Unlike pragmatic actions, they do not reduce the distance to the goal. Instead, they

[†]Although McCarthy and Hayes (1969) originally posed the frame problem in terms of a specific representational framework, the situation calculus, the frame problem as more broadly construed concerns a system designer's problem of finding a representational form that eliminates the need for the system to reason about things that tend not to change (see Janlert 1996).

[‡]In Tetris one of seven shapes, each composed of four squares, descends the playing area. The object is to fill a row by positioning the shape through rotations and translations. Once filled, the row disappears. The shapes descend with increasing speed, and the game ends when they fill to the top of the game area (see <http://en.wikipedia.org/wiki/Tetris>).

reduce or simplify computation by offloading mental tasks to the environment. For example, rotating a game piece physically can save the effort of rotating it mentally and simplify the mental task of matching the piece to a contour. Physical rotations can even be used to identify a piece before it has fully appeared. These results resonate well with work on active vision in robotics, in which perceptual-motor routines reduce the reliance on internal representations by computing information on demand (Ballard 1991). When discussing epistemic action, it is important to note that earlier stages of mental processing can activate motor responses directly. Thus, the delay between sensing and action can vary from as low as a few milliseconds to seconds or even longer depending on the kind of processing involved.

Since epistemic actions enable people to *alter their cognitive states* instead of merely changing the physical state of the world, they are often used socially. Thus, people seem to use gaze in ways that are often epistemic. When they do so, like other action, the effects are visible to others. They make public (at least in part) what someone else is thinking and feeling. For example, the way in which a person breaks eye contact may indicate that the person is thinking hard about an answer. In other circumstances, it may indicate that the person is feeling self-conscious—or even nervous about the prospect of telling a lie. MacDorman *et al.* (2005) found that the downward gaze tendency of study participants who were contemplating an answer when interacting with a human questioner or an android they believed to be under human control disappeared when they believed the android to be acting autonomously. What the participants believed about the android seemed to influence whether they showed modesty with their eyes. So clearly in the social domain, epistemic action does much more than simplify mental tasks by performing them in the environment. This is because it exhibits intentions, empathy, and consideration, and thus prompts other people to act as a ‘cognitive resource’ (cf. Clark 1997, Lee 2006). The action ceases to be centred on the individual and, instead, becomes a means of co-ordinating interaction partners. It is a paradigmatic example of *distributed cognition* (Hutchins 1995, Clark 1997, Alac 2005).

Yet the significance of any particular action cannot be evaluated without some kind of standard against which to measure it. *Interpersonal norms provide evaluative standards in terms of which observable behaviour can be interpreted.* Conformance to norms and deviations from them have value from a given person’s perspective (Christensen and Bickhard 2002). People orient to norms and exploit each other’s orienting. In mother–infant dyads, joint events stabilize around norms as the baby’s brain self-organizes under the dual control of mother and infant (Trevarthen 1998). The affective dynamics of their relationship help to shape this process. Although norms develop based on their selection history, they are understood as performing a function for a broader system or systems (Davies 2001). Norms may be characterized at many levels of description (Hinde 1987) and organization, ranging from an individual’s quirks and idiosyncrasies, to the ways of relating that develop in intimate relationships, to group and inter-group practices, to societal customs. Although in philosophy norms are identified with sentences that prescribe how things ought to be, they may be construed more broadly to include, for example, meaningful ways of interpreting patterns that emerge from closely co-ordinated interaction. If we are to build androids capable of being transformed by social or cultural organization, they must be attuned to norms. Like baboons, babies and Tetris players, relationship-sustaining machines need to fit signals to social constraints. Because it is likely that humans would *experience* such androids as intelligent, this paper explores how, in two species, norms support modes of action.

1.1 Outline

This paper goes beyond Kirsh and Maglio’s (1992, 1994) seminal exploration of epistemic and pragmatic action to suggest that the ‘human interface’ couples two corresponding perceptual

mechanisms[†] with affect and the use of words. Even if language is ignored, norms function both within and beyond individuals, serving individual, interpersonal, group and cultural-societal functions. Then, in comparing primates, while pragmatic action dominates the baboon's world, a caregiver's epistemic actions shape the human baby's world from birth. Long before infants learn to talk, affect and the caregiver's use of cultural norms influence their decision-making. As in Tetris, these ways of co-ordinating activity lead to changes over time. In considering machines, there is no reason, in principle, for android decision-making not to use similar movements. Finally, engineers are urged to design androids that use the 'value' of patterns in such a way that robots, like Tetris players, will be able to couple modes of action that are both epistemic and pragmatic. Under these circumstances human response can be used to soft-assemble sophisticated human–android co-action.

2. Epistemic norms at the *human* interface

While the advent of the computer age spawned many theories comparing computers to human brains (Newell 1980, Pylyshyn 1984, Johnson-Laird 1988, Dennett 1991a), as computers merged into the social background, they ceased to amaze. Although they have transformed many of our social structures, a reliance on rules and formal tokens excludes the use of norm-based patterns. Not unlike wax-tablets or printed products, their powers depend on how *we* use notations. The intellectual excitement of comparing brains to computers now seems to say more about human imagination than machines. The problem is that current designs preclude the effective use of computers in situated events:

'Computers have no choice about when or what to produce. They do not have the option of breaking rules . . . They cannot integrate their production with the (social and situational) circumstances in which they are made to produce.' (Taylor 1997: 150).

Of course, computers have been programmed to model norms in the narrow sense of statistical regularities; this has been a major focus of research on learning in artificial neural networks (Bishop 1995). A number of robotic systems have been built in which robots affectively appraise their environment based on internal drives directed toward maintaining physiological or emotional equilibrium (MacDorman *et al.* 2000, Breazeal 2002, 2003, Arkin *et al.* 2003, Sawada *et al.* 2004). The behaviour of these robots is normative in so far as it functions to maintain the well-being of the whole 'organism'.

What so far has been missing is an understanding of *how robots could dynamically contextualize their activity during social interaction*. Context cannot be predefined; it emerges from the interaction (Dourish 2004). However, contextualization requires an appreciation of interpersonal and cultural norms. In principle, there is no reason why 'sociable robots' could not be extended to pick up on these kinds of cultural norms, which babies have already started to learn at 3 months (Cowley 2003, Cowley *et al.* 2004). Even in Tetris, human action depends on situated events. We extend our neural powers by using a screen for decision-making. Instead of slotting a piece home, we may choose to rotate it or shift it sideways. Far from fitting *rules*, epistemic decision-making draws on experience in dealing with circumstances. Must android actions be rule-bound or could their decision-making be more like that of human Tetris players?

In playing Tetris, we often do *not* act in goal-oriented ways. To contrast humans with machines, Kirsh and Maglio wrote a program RoboTetris to implement the following:

[†]The perceptual mechanism to support epistemic action recognizes a pattern worthy of epistemic action, while the perceptual mechanism to support pragmatic action recognizes the meaning of the pattern in relation to the goal ('This piece can be fitted here').

- (1) Create an early, bitmap representation of features of the current situation.
- (2) Encode the bitmap in a more compact, chunked symbolic representation.
- (3) Compute the best place to put the piece.
- (4) Compute the trajectory of moves needed to achieve the desired placement.

Computers running RoboTetris outperform humans. Every action that they take is pragmatic: it is taken to position a piece. Just as Taylor suggests, the computer does not make any choices about when or what to produce. Its action does not allow any rule breaking, and there is no need to integrate experience with action. The importance of Kirsh and Maglio's (1994) study is to clarify how this differs from human play.

Expert human players exhibit (information-seeking) epistemic action 'in abundance almost from the moment a piece enters the screen' (Kirsh and Maglio 1994: 524). This phenomenon could be established, because every touch of the keyboard can be measured against the current state of play. Experts often make shifts and rotations that, from a programmer's perspective, are redundant, occurring even before the player has decided what to do. Strikingly, experts use complex norms in deciding when to use rotation. For example, they are *more* likely to rotate pieces with double ambiguities (shape-and-position versus position) or ones whose visible ambiguity will only be resolved slowly. At times, moreover, pieces are rotated *before* their shape is fully visible. How can this be explained?

Given a human body, a task can often be simplified by drawing on perceptual-motor routines. Instead of making a judgement call about piece placement, the player can ensure a piece's correct alignment by shifting it to the side and counting key taps while shifting it back. Mental rotations can often be avoided when norms prompt us to move a piece on-screen in ways allowing us to see where it fits. In Tetris, epistemic actions are a micro-response to 'What now?' In 'acting before thinking', we simplify the game. For Kirsh and Maglio, the discovery of this type of action has major implications for cognitive science.[†] Among other things, folk views are seen as fundamentally wrong: intentions often do *not* drive voluntary actions. Unlike RoboTetris, humans manipulate the world to enhance their mental performance. The computer game itself simplifies their cognitive tasks. In epistemic action, therefore, *we act to alter our cognitive states*. Especially in time-windows of a few hundred milliseconds, we often avoid complex cognitive processing by using experience to do something. Indeed, as Kirsh and Maglio emphasize, epistemic actions are at the fringe of awareness (Mangan 2001) and, for this reason, are rarely reported in telling others about playing the game.

Kirsh and Maglio's work is highly instructive on the human 'interface'. For example, they observed that, as people grapple with Tetris, their decision-making process makes use of *norms derived from the interplay of affective states and experience*. Since these arise while interacting with the program, Tetris features complex decision-making that integrates pragmatic and epistemic actions during real-time play. While spoken instructions to the novice may initially motivate pragmatic moves, epistemic moves are carried out on a much simpler level and depend in part on some form of operant conditioning. Furthermore, interviews show that people use a third level of organization. Although this perception is generally false, players believe that actions depend on 'rules of thumb' (e.g. 'I stack pieces up in the middle while waiting for a long piece').[‡] Human decision-making thus spreads over three timeframes. Epistemic action is measured in milliseconds: we rotate pieces as they become visible. Pragmatic action uses what we see in the timeframe of a few eye fixations. Finally, at the macro level, we alter our

[†]While Kirsh and Maglio are modest, the nearest concept to epistemic action may arise in modelling values as serving a kind of goal (Campbell *et al.* 2002). Whereas Kirsh and Maglio see the 'goal' as that of *changing one's* state, they present the change as fulfilling a (pre-set) goal.

[‡]As so often with folk-psychology, explicit beliefs may be useful without being any guide to causal mechanisms; because they are normative, they can serve both action and perception.

actions in ways that impact future epistemic and pragmatic action. Our ‘policy’ moves play out over seconds or minutes (Clark 1997). They are, of course, inseparable from language.

Both epistemic and pragmatic actions use pieces as signals. However, epistemic action is simpler and involves shorter timescales, because it is based on short cuts in human information processing, for example, from perception to motor control, thus avoiding cognitive deliberation. So while pragmatic action depends on sophisticated events including the ‘mental’ rotations, matching and plans implemented in RoboTetris, epistemic action exploits rudimentary cues. It allows us to compensate for bodily and cognitive limits by setting ourselves up for what comes next. We develop modes of action based on a network of affect-coloured events that prompt learning about the effects of shifting and rotating pieces. Context-sensitive moves lead to a rise in useful (and useless) habits. In spite of this, epistemic action is on average beneficial, producing useful cognitive changes. Rotations and shifts improve performance even if, from a competence perspective, they are redundant. This is why epistemic action becomes more prominent in expert play. Equally, a player experiences epistemic moves as purposeful and evaluates them. The moves show that a person’s ability to play Tetris depends on sub-personal control systems.

In rotations, the screen acts as a representational arena. Experience enables current concerns to be integrated with actions that use affect-marked prompts. Going beyond the information given, norms prompt us to action. An expert senses, for example, that *this piece now* is emerging and ambiguous. Seeing it may thus prompt a routine of rotate-and-look. Ambiguity may trigger this action, which is not deliberate. Indeed, a decision to act epistemically may correlate with both the current state of play and how the player feels. Such modes of action develop from affect-based perception. Epistemic action arises without needing processes as complex as perception–action matching or rules of thumb. Often epistemic action draws only lightly on awareness as it uses causal process between brain, body and world. Human reliance on epistemic movements in Tetris suggests that, in coming to terms with an indeterminate, perceptually ambiguous world, they are of even greater importance (see Ballard 1991). Given their benefits, could androids be designed and trained to carry out something similar?

3. A biosocial view of norms

As epistemic actions are performed in Tetris, what counts is *when* we choose to shift or rotate a piece. Since the answer to this question depends on norms, it is necessary to examine their nature more closely. Norms saturate the biosocial world; however, many theorists contrast the causal with the normative (see Christensen and Bickhard 2002). Furthermore, because norms are usually dealt with in social theory (see Therborn 2002), nobody has systematically connected the norms of nature to their social, institutional and cultural counterparts. This disconnection is exacerbated because, for many purposes, it is enough to treat norms as abstract categories that are irreducible to physics.[†] In obeying laws or following habits, for example, we need not know about the grounds for our behaviour. Equally, in Tetris, an epistemic action is a real-time response to the affective value of a physical complex.

Androids must register and orient to norms to sustain relationships with humans. This is because norms are ubiquitous in the world of human interaction. Thus, Goffman’s (1963) work on the physical events of human encounters shows us a world organized by ‘moral’ norms. Co-presence is dominated, not by words, but by microsocial events that arise as we co-ordinate our activity (Cowley and MacDorman 1995). In the ‘interaction order’ (Goffman 1983), we

[†] Conte and Castelfranchi (1995) take the view that machines can ‘adopt’ norms. Rejecting any top-down approach, we argue that machines should sensitize to norms that, while irreducible to physics, use real patterns (Dennett 1991b).

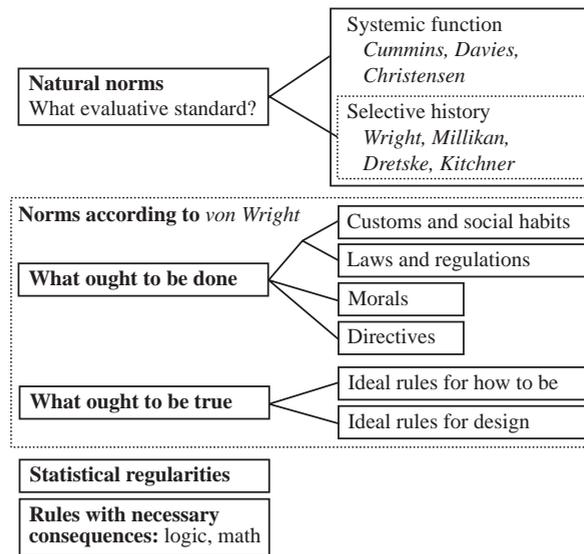


Figure 1. An incomplete taxonomy of the meaning of norms. Davies (2001) conceives of natural norms as evaluative standards concerning the contribution of a process or mechanism to a higher-level system. (The higher-level system could be the population with natural selection as its causal process.) Von Wright (1963) ignored natural norms and took a third-person view of cultural and deontic norms and ideals. This conceals the distinction between *norms*, which multiply constrain an organism in its environment, and *norm formulations*. A construction sign flashing 30 literally tells drivers not to exceed 30, but the actual effect this norm formulation will have depends on many factors, including road conditions, the flow of traffic, and so on. Individuals orient to the higher-order system of driver-in-traffic. Simply to follow the sign's prescription could cause an accident!

do not use explicit directives but rather the rule 'fit in!' (and when we choose to stand out, norms let us know what we are doing). Gradually, we discover 'modes of action' that can be used to define our goals (Goffman 1983: 8). Even 'How are you?' may be said with sensitivity to who is being addressed, circumstances, and affective display. Norms are so dominant that, as von Wright (1963) noted, the term has an unusual ambiguity. This is shown figure 1.

While relationships depend on how human beings use social norms, similar organization occurs throughout biology. For our purposes biological systems are treated as paradigmatic (Davies 2001). This is useful because, unlike their interactional and social counterparts, the norms of nature clearly have a physical grounding. Davies (2001) used work by Wright (1973, 1976) and Cummins (1975) to show why. Although enacted by physical processes, norms within organisms produce functions that are, in many cases, extended by recurrent effects and selection mechanisms. With Cummins, then, the norms of nature can be defined as 'causal or structural capacities that contribute to the exercise of some larger systemic capacity within the larger systemic view' (Davies 2001: 4). A classic example is the heart's functioning: this norm permits the higher-level operation of ensuring blood circulation.[†]

Although it is clear that biological systems need norms, Davies's theory has a broader application. As primates, we use norms that are based in biology *and* on interaction with the physical and social environment. In showing how to bridge this divide, Davies unwittingly makes a substantial addition to systemic theory. By requiring norms to cross hierarchical

[†]Davies (2001: 53) points out that, apart from pumping blood, the heart also functions to ensure the circulation of nutrients, to increase reproductive output and, under certain conditions, to contribute to the distribution of genotypes in the population. Systemic capacities typically develop layers by embedding the proximate functions of a trait in higher levels of organization, including biological populations and cultural groups. Thus, Davies argues that even Wright's (1973, 1976) selective account of norms can be understood as systemic when selection is viewed as a causal process that maintains the presence of a given feature within a changing population system.

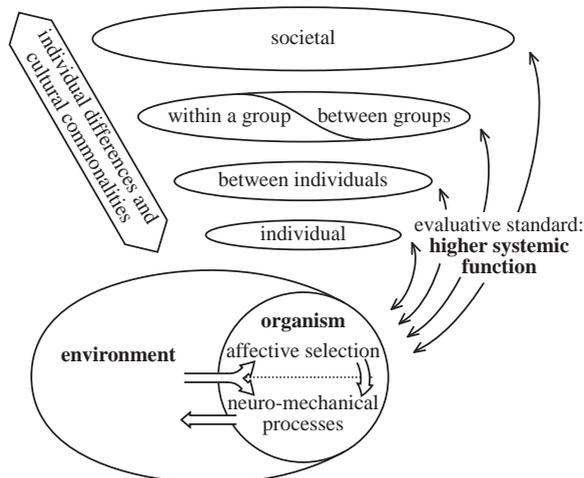


Figure 2. Unlike von Wright (1963), we take a first and second-person view of norms: What counts as a norm for an actor and co-actor in a relationship? Our claim is that, like the norms of nature (Davies 2001), interpersonal norms can be characterized in terms of their contribution to the function of higher-level systems (e.g. the mother–infant dyad). At different levels (e.g. individual, interpersonal, group, societal; see Hinde 1987), these norms both derive from their systemic functions and regulate systemic activities. Our focus is on what Goffman (1983) called the *interaction order*: a world of norms to which participants in social life are bound to orient. Norms, however, reflect individual differences. The norms of epistemic action in Tetris differ among players, reflecting their individual histories, bodies and cognitive capacities.

levels by contributing to the function of a larger system, Davies not only distinguishes them from ‘dispositions’ but also abolishes the need for top-down mechanisms (see Christensen 2004, Wheeler 2005). According to this view, epistemic action can link a lower-level feeling of what happens with higher-level causal events. Thus, playing Tetris is norm-based in that it links the individual system to the higher-level unit of player-in-the-game.[†] For those interested in modelling behaviour, the advantage of invoking the effects of interacting systems is that this enables them to conceptualize how the causal events of biology become enmeshed with social life. *Through bottom-up processes, norms enable individuals to engage with recurrent patterns that constitute higher levels of organization beyond the body* (see figure 2). This distinguishes norms beyond the body from their more biological counterparts, although both can serve functions and act as evaluative standards. While this applies to other primates, it is especially important for language users. In our species, judgements use both actions and words. In Tetris, for example, the human interface prompts simultaneous use not only of objectives and affect-marked patterns, but also of rules of thumb or verbally formulated beliefs.

A similar logic applies to primate social lives. Just as evaluative standards can benefit a norm user, they allow conspecifics to undertake social assessment. Accordingly, when social norms become established, they produce control systems in the interests of the group. In animals like baboons and human beings, therefore, pure ‘behaviour’ falls away. Actions are so multiply constrained that *any* move or failure to move provokes norm-based judgements. Where language is involved, these judgements can also be analysed against how we act and affect social attitudes. As Blackburn (1999: 51) remarked, ‘we do not know what someone is doing until we can see their mere movements as expressions of purpose and intention’. The importance of this for android science can hardly be exaggerated. To build relationships,

[†] Although beyond the scope of this paper, rudimentary perception of, say, something in row five triggers a low-level system akin to core-consciousness (see, Damasio 1999) to interact with high-level affect-saturated systems. A player may report, say, a sinking feeling, or thinking ‘oh, blast!’

we must experience the machines as taking our movements as expressions of purpose and intention. However, this does *not* imply they need to ascribe sense to behaviour. Like babies and baboons, it may suffice if they can treat adults as producing movements with affective values. Pursuing this, let us consider how interactional norms contribute to non-linguistic members of two primate species.

4. Interactive learning in baboons and human infants

In Tetris, pragmatic and epistemic actions develop as skills and are honed against reliable causal processes. Players orient to norms so that, unawares, they draw on past successes and failures. Having considered how we learn to play Tetris, let us turn to social events. Here we orient to multiply constrained systems that draw on social and biological norms. Owing to the complexity of encounters, when we examine how baboons and babies use social norms in the ‘strategic signals’ of primate co-ordination games, pragmatic and epistemic action can no longer be so neatly distinguished (Ross and Dumouchel 2004).

Hamadryas baboons are, in many ways, like us. They live in male-dominated groups where energy goes towards fighting, feeding and sex, while social life is ordered around intricate norms. This gives rise to flexible, adaptive behaviour, which, when viewed closely, shows species-specific characteristics. In place of language, Kummer (1995) stressed, they use gestures, of which 98% achieve identical effects in Arabia, Ethiopia and at the Zurich Zoo. Their co-ordination games are such that ‘No baboon can learn from another a gesture that is not programmed by it hereditary material’ (Kummer 1995: 29). Despite hard-wired signals and effects, baboon behaviour is far from automatic. Like Tetris players, they expend limited cognitive resources as social norms allow conspecific movements to be experienced as expressions of purpose and intention.[†]

In illustration, consider the social triangles that dominate hamadryas baboon life. In some ways, the three-way manoeuvring has familiar features. Infants deal with individuals by using norms and social signals to build relationships and, using resources beyond the body, place themselves in the interaction order. In spite of parallels, the social triangles also play out in ways we find strange. For example, Kummer (1995: 39) describes an occasion where Wiffel, a youngster, play fights with an older male. During the game, Wiffel is hurt and, as *always* with hamadryas baboons, moves towards the largest present conspecific. To our surprise, instead of punishing Wiffel’s assailant, this male mounts a frightening sham-attack on the already screaming Wiffel. These events reflect on norms, which an observer might formulate as: (a) if threatened, move towards a large male; and (b) if a crying youngster approaches, show aggression towards something. Indeed, hamadryas ‘protectors’ threaten refuge-seeking baboons just as often as they do their assailants. Unlike a human being, the baboon does not seek sympathy: it does not use the world to alter its cognitive state. Rather, it acts in the way that Kirsh and Maglio deem *pragmatic*: it acts to get protection and only indirectly regains ‘its own inner equilibrium’ (Kummer 1995: 39). Acting pragmatically to regain equilibrium has been the approach taken by social robotics (Breazeal 2002, 2003, Arkin *et al.* 2003, Sawada *et al.* 2004).

Hamadryas baboons seem more skilled in assessing behaviour than in managing others. While epistemic action may occur, their doings seem designed to alter the physical world.

[†]Although baboons ‘know’ in some sense what they are doing, this ‘knowing’ cannot be formulated and is thus intrinsic to norm-based action. In that evaluative standards cannot be separated from the behaviour-in-context, the events are not ‘normative’.

Why should this be? Since baboons are good at social learning, compared to us they show a pragmatic bias.[†] For example, we read of the Zurich Zoo baboons:

'There was a special, mild form of bite that the adults in the group used in dealing with troublesome children. The young one was grasped by the forearm or foot, which was then raised to the adult mouth and then gently bitten It never made them scream.' (Kummer 1995: 33).

Kummer (1995: 33) sees this 'kindly admonition' as a group invention based on social learning. It may be partly epistemic because even if only as a side-effect, the mild bite can alter the adult's cognitive state of being troubled. But it seems to be less cognitive than motivational. Despite their skill at social learning, baboons exert 'limited choice' about when and what to produce and, strikingly, do not 'break rules', and thus bear a resemblance to Taylor's (1997) description of rule-bound computers. While 'using social and situational circumstances', they rarely reintegrate outcomes in these circumstances with their own cognitive states.

The 'pre-moral' function of social triangles provides another example of this pattern. Hamadryas baboons are polygamous: *some* males control two or more females, while many young males remain unattached. A number of experiments investigated the underlying norms. For example, when a female is introduced to two males, she invariably bonds with the dominant one and, from the moment of bonding, the subordinate loses all interest in the female. To test the hypothesis that this behaviour is based on motivational states, bonded males were captured and separated from their females. Remarkably, other males then bonded with the 'wives', even though the dominant male could see them. Clearly, the norms that prevent unattached males from showing interest in females differ from moral rules! On release the dominant male usually got his 'wives' back, but never punished them or the abductors. For Kummer (1995), the hamadryas convention 'solves the problem of competition without force'. Since baboons do not strategize, pre-morality need not be formulated or taught. This suggests that hamadryas signals or norms are rarely used in self-serving ways and that, unlike humans, these primates do not use the world to change their cognitive state.[‡]

In such respects, baboon social intelligence contrasts both with that of Tetris players and human infants, because both have the ability to provoke epistemic effects. For example, newborn babies carry out mimetic actions that use adult gestures to respond in kind (e.g. Melzoff and Moore 1997, Kugiumutzakis 1999). When an adult sticks out the tongue, for example, an infant may reciprocate, which is almost invariably rewarded by an affective response. From the start, moreover, this will not be purely emotional, but bound up with cultural norms, that is, evaluative standards. In Japan, the example may embarrass the audience, while in South Africa the same story sets off happy laughter; because of the pervasiveness of norms, even tongue-showing produces epistemic effects. Just as in Tetris, we are attracted to using predictable events by drawing on affect-marked response. In this case, however, infants line their behaviour up with cultural values. Unlike baboons, they develop interpersonal routines that, using local norms, alter their cognitive states. From birth, they are subject to displays of norm-based response or 'context-altering behaviour' and, by 3 months, use these in co-ordination.

[†]Marais (1939) made striking observations of social learning in South African *cachma* baboons. For example, local baboons adopted the practice of tearing open the entrails of lambs to eat the curdled milk. However, when food became scarce, they developed a new method: 'The lamb is no longer killed by tearing open its entrails: it is stretched out on its back, both jugular veins are carefully exposed and severed by biting, and then the animal is left until it is dead and the veins are emptied of blood, for it seems that the baboon still has a dislike of blood. Thereafter the lamb is properly skinned (in a manner in which no beast of prey could do it) and the flesh torn from the limbs and eaten' (Marais 1939: 11–12). The complexity of this procedure illustrates baboon intelligence and, as the method was shared, social learning. As with the hamadryas baboons, however, adaptive, flexible intelligence seems to be directed at altering the physical environment.

[‡]Another account would appeal to 'theory of mind'. In comparison, focus on epistemic action has the advantage of parsimony. It allows cultural organization to produce behaviour patterns between the body and its world. In baboons, practices link a physical environment to 'conventions'. The apparent lack of epistemic action may reflect a world largely without arbitrary practices.

Remarkably, this is before a baby can carry out even simple pragmatic actions like reaching for an object that it wants to suck. In our species, uniquely, epistemic action precedes its pragmatic counterpart.

In the first 6 weeks, human infants are generally concerned with adapting their bodies to the physical world. Unlike in baboons, interaction is mainly reactive (Rochat 2000). Then, before even starting to act on the physical environment, the baby gets caught up in epistemic events. In the first infant ‘revolution’, after 6 weeks, the infant discovers co-action (Stern 1977, Trevarthen 1979, 1998, Kaye 1982).[†] The infant’s social and affective motivations come to show the stamp of the caregiver’s world (Trevarthen 1988). Already, the infant holds attention and acts to provoke culturally salient effects, evoking reactions that exert predictable effects on its own cognitive state. Even small infants meet, flout and, at times, surpass norm-based expectations by, for example, responding to culturally specific signals to fall silent (Cowley *et al.* 2004). Not only is co-action managed around affect, but also, during these events, parents make constant evaluations (Stern 1977, Cowley 2004). Babies develop a systematic capacity to use historically based experience and, in so doing, develop norms that use caregiver expectations. While based in mechanisms for closely timed control over movements, the child’s doings are treated as actions. Using operant conditioning, the infant begins to act in ways that ‘work’ by changing its cognitive state. Infants treat their caregivers as people (Legerstee 2005). While limited in when and what to choose, infants also break the rules: in response to signals for silence, for example, they may continue crying. Although like baboons in developing a capacity to integrate understanding with social and situational circumstances, they differ from these primates in that 3-month-old human infants already use expression and even vocalization to social ends (Cowley *et al.* 2004, Cowley 2004). It looks as if epistemic action is intrinsic to human nature. We tend to approach the world and other people as cognitive resources.

In some ways, the infant world is like a Tetris screen. Just as Tetris pieces prompt players to develop norm-based routines that scaffold on co-action with the game, using cultural norms, other people prompt infants to develop norm-based routines that scaffold on their co-action. Since infants also react to the reaction of others, they develop hierarchical levels of self-control. The early phase dominated by epistemic action gives way, however, to a stage dominated by motor development (cf. Lindblom and Ziemke 2006). In the period from 4 to 9 months, while focused on pragmatic action, babies still use interpersonal routines that promote fun, as the situation requires. Then, at the end of the first year, a second ‘revolution’ occurs. The baby discovers ways of meshing its capacities for epistemic action with skills based in pragmatic action. This gives rise to triadic behaviour that seems to be exclusively human: talk and play come to resemble tickling and teasing as well as primate signalling.

Human infants differ from baboons in the following ways:

- (1) Infants are born in a pre-pragmatic developmental state.
- (2) Infants live in a world where high-level hierarchies are cultural.
- (3) Human caregivers use norms that connect with the culture level.
- (4) Social learning and operant conditioning allow high-level norms to prompt an infant to become a person-in-a-relationship.

From birth, the caregiver’s verbal formulations educate the child, not only with respect to language ability, but also, more importantly, by demonstrating how the norms of language shape adult modes of action. In a sense, a baby is already engaged in language. Without

[†]Trevarthen calls this *intersubjectivity* on the grounds that it depends on an intrinsic motivational system that links people as subjects. Others invoke dynamical capacities that prompt a baby to develop expectations (Stern 1977, Kaye 1982). By referring to co-action, we avoid the debate.

needing a special neural system (Arbib 2002, Hauser *et al.* 2002, Tomasello 2005), decision-making enables the epistemic to mesh with the pragmatic. Understanding uses a motivational system where affect prompts action. The child builds a ‘human interface’ around the norms that allow it to enact human relationships. Not only does this process ensure that, by 10 months, infants have developed attachment styles, but also, by responding as they do, adults prompt babies to orient to what they value. For the adults, of course, these norms can also be explicitly formulated in ways that allow social learning to depend on co-action. In short, using cultural norms is helpful in seeking to exert control over a baby’s behaviour.

5. Human–android co-action

From an engineering perspective, we face two questions. First, can androids mimic something like primate intelligence in orienting to both signals and social constraints? Second, can robots make use of epistemic action? Epistemic actions are more rapid than their pragmatic counterparts mainly because they are less computationally costly, relying on more rudimentary perception and movement. What an engineer can take from human babies and Tetris players is that, while they effortlessly link the causal and the normative, they use simple mechanisms. Baboons do not seem to use comparable resources, because their signals are hard-wired. Rather, they need extensive experience of the world so that gradually they gain control over motivational states. Must androids, then, depend on baboon-style hardwiring? Or can they, like infants, augment their intelligence by using humans as cognitive resources? Can they be designed so that their control systems develop from using expressions of purpose, intention and value?

Pragmatic actions are central to complex behaviour. As we have seen, their power can be augmented not just by ‘knowledge’, but also by affect-based interpersonal routines. If androids can develop something like epistemic action, they may indeed sustain relationships with people (MacDorman and Cowley 2006). However, they must first elicit the right kinds of responses before they can use sophisticated timings to react to these responses in ways we find ‘transparent’. They must also use experience to anticipate our reactions by acting in ways that strike us as reasonable. With this in mind, let us consider human response to interaction-oriented robots. Even mechanical-looking humanoid robots elicit norm-based response. For example, in a Japanese primary school, the robot Robovie elicited positive response and, for some children, became a valued companion. In expressing what Cowley and Kanda (2005) call ‘partnership’, the children used norm-based behaviour and, indeed, imagination. For example, when the robot used pseudo-learning to ‘confide a secret’, this was heard as deliberate action. Responding, a child did a dance of delight before sharing the news with her classmates and rewarding the machine with manifest joy.

Not only do robots elicit norm-based responses, but responses can at times also be predicted in advance. It is imaginable, therefore, that robots could be pre-programmed to respond ‘transparently’—at least in shorter-term interactions that do not involve mutual adaptation. More dramatic effects can occur with androids: even those who build androids orient to culturally based norms. For example, at Hiroshi Ishiguro’s laboratory at Osaka University, engineers enjoy ‘teasing’ the female android in front of visitors. They do so, for example, by touching her shoulder so that she moves to fend them off. The effect depends, of course, on local norms pertaining to when and how male engineers are permitted to touch women. We also respond to machines using norms associated with eye contact (MacDorman *et al.* 2005). These examples show, moreover, that by designing androids that look and move like us, we can make their movements seem transparent. The challenge lies not in getting humans to react naturally

(Turkle *et al.* 2006), but in ensuring that androids use *aspects* of these reactions so that their responses strike us as purposeful and intentional (M. B. Cowley 2006).

The engineer needs to get machines to identify and respond transparently to context-changing behaviour. Such response can carry culturally specific values and, in some way, be appropriate. The machine will identify likely norm-based behaviour and, in real-time, movements that allow for evaluation of how the human can ‘go on’ in the interaction (Wittgenstein 1953). From an engineering point of view, machines need biases that facilitate response to epistemic action and, conversely, to use such actions to set up their use of motor routines. As with humans, they will use reactions that we understand to anticipate how people will respond. Internally, then, such machines need to have an intrinsic motivation system that allows use of affective signals that give weightings to ways of acting and, especially, actions that are ‘meant’ to change their state.

Let us imagine a robot that responds to a child’s dance of delight. For the sake of argument, it might use timing to identify visible and audible excitement, check its positive feel and, using an interpersonal routine, select a canned utterance of ‘This is our secret’. Such a robot would use a human-looking ‘android interface’:

- The robot perceives that, at a given time, behaviour is likely to have a given value.
- It identifies a pattern and, using certain physical characteristics, selects a response fixed by its value.
- The value-weighted response is transparent to humans.
- Subsequent human response is also assigned a value.
- This is used to give an appropriate reinforcing reaction.

Particularly if the pattern and the biases used in ascribing a value to the pattern are grounded in experience with an individual, the machine begins to simulate a (language-deprived) person. Using sensitivity to norms, it begins to assess response. Finally, it can develop norm-based social learning.

6. Whole agents in a cultural world

Turing (1950) recognized that the experience of a human subject would vindicate tests of intelligence. Particularly if androids engage us in co-action, we edge closer to the goal. Indeed, even mimicking baboon social intelligence would mark progress in understanding why, during social behaviour, we experience some movements as intelligent. Merely by emphasizing affect and timing, we begin to escape from the focus on symbolic language processing that had dominated much work in artificial intelligence. What matters in whole agents is not how we analyse behaviour symbolically, but how flexibility enhances adaptation. Organisms use multiple constraints to develop ways of acting that benefit individuals and groups. In so doing, decision-making recruits norms based both within and beyond the body.

It is of intrinsic interest to design machines whose inner systems deal with perceived patterns in the social environment. While these may mimic baboon-type interaction, relationships are at least as dependent on epistemic action. In both cases, however, decision-making must deal with both experience and events on at least two timescales. Even baboon-like social interaction uses perception and action that are guided by past encounters with particular individuals. Without this ability, androids cannot perceive an individual’s movements as actions, use norms to deal with signals, or develop interpersonal relationships. Conversely, if their bodies can be used to develop a decision-making interface, we can challenge Taylor’s objections to machine intelligence. In contrast to computers, androids that co-act with humans will be felt to ‘make

choices', 'break rules' and reintegrate actions with changing perceptions of the world. Like primates, they will use imitation while using social norms.

Primates use conspecifics to extend their cognitive powers. Much of what we do depends on movements that adults perceive against complex sets of norms. In humans, however, it seems that these are extensively used in affecting our cognitive states. Thus, as Goffman (1955) noticed, almost anything can make a contribution to *face* (i.e. self-image as defined by approved social attributes). Considering how such responses give advantage to infants brings out the importance of the observation. First, as Deacon (1997) argued, the competitive processes that shape neural organization enable environmental stimuli to impact on underdeveloped systems. Second, since we perceive infants as cute (and at times horrible), appearance and basic behaviour elicit affectively charged response. Third, owing to operant conditioning, specific movements can be time-locked with an affective value. Fourth, while not well understood, this allows a caregiver to sense what is likely to happen. Fifth, caregiver 'thoughts' mediate co-action. Using cultural norms, caregiver beliefs give actions predictability associated with real-time expressions of affect. Indeed, changing motivations may underpin the events that, in a few months, transform an infant into a child. Much, therefore, can be gained from motivational systems that respond during co-action. As promised by MacDorman and Ishiguro (2004), androids of this design may indeed sustain relationships. Such machines, however, need not be modelled on babies: rather, the key lies in designing them so that their causal processes can co-opt our sensitivity to social norms. The analogy, then, is with the developing skills of Tetris players. Like them, androids can learn to exploit affect-saturated signalling to cognitive ends.

In designing social machines from the bottom up, much can be gained by exploring how norms connect hierarchical levels. At some point, however, epistemic action must not only complement but also augment baboon-like social intelligence. Machines must come to use the non-social world, and, of course, signals that are independent of affect. Eventually, therefore, an android will need an interface with a third level that, like a Tetris player's, allows it to formulate rules of thumb. Today, however, the challenge is to develop 'mindful' machines that use movements that can be experienced as expressions of purpose and intention. If epistemic action is mimicked, like Tetris players, the world can be used to train an android's behavioural and perceptual systems. This could be a move toward freeing the social and behavioural sciences from folk psychology. Accordingly, relationships use experience of norm-based co-action where, using affect, each party serves the other as a cognitive resource. Given a suitable interface, androids can soft-assemble *modes of action* that function as more than movements. Alongside control systems that use us as cognitive resources, human–android co-action can take place in a world that, like a Tetris screen, promotes the rise of systems that prompt norm-based judgements.

References

- M. Alac, "From trash to treasure: learning about brain images through multimodality", *Semiotica*, 156, pp. 177–202, 2005.
- M. Arbib, "Grounding the mirror system hypothesis for the evolution of the language-ready brain", in *Simulating the Evolution of Language*, A. Cangelosi and D. Parisi, Eds, London: Springer, 2002, pp. 229–254.
- R. Arkin, M. Fujita, T. Takagi and R. Hasegawa, "An ethological and emotional basis for human–robot interaction", *Robot. Auton. Syst.*, 42, pp. 191–201, March 2003.
- D.H. Ballard, "Animate vision", *Artif. Intell.*, 48, pp. 57–86, 1991.
- M.H. Bickhard and L. Terveen, *Foundational Issues in Artificial Intelligence and Cognitive Science: Impasse and Solution*, Amsterdam: Elsevier, 1995.
- C.M. Bishop, *Neural Networks for Pattern Recognition*, New York: Oxford University Press, 1995.
- S. Blackburn, *Ruling Passions*, Oxford: Clarendon Press, 1999.
- C. Breazeal, *Designing Sociable Robots*, Cambridge, MA: MIT Press, 2002.

- C. Breazeal, "Emotion and sociable humanoid robots", *Int. J. Hum. Comput. Stud.*, 59, pp. 119–155, 2003.
- R. Brooks, "Intelligence without representation", *Artif. Intell.*, 47, pp. 139–159, 1991.
- R. Campbell, J. Christopher and M. Bickhard, "Self and values: an interactionist foundation for moral development", *Theory and Psychol.*, 12, pp. 795–823, 2002.
- W. Christensen, "Self-directedness: a process approach to cognition", *Axiomathes*, 14, pp. 171–189, 2004.
- W. Christensen and M. Bickhard, "The process dynamics of normative function", *Monist*, 85, pp. 3–28, 2002.
- A. Clark, *Being There: Putting Brain, Body, and World Together Again*, Cambridge, MA: MIT Press, 1997.
- R. Conte and C. Castelfranchi, *Cognitive and Social Action*, London: University College London Press, 1995.
- M.B. Cowley, "The relevance of intent to human–android strategic interaction and artificial consciousness", in *Proceedings of the 15th IEEE International Symposium on Robot and Human Interactive Communication*, 2006.
- S.J. Cowley, "Distributed cognition at three months: mother–infant dyads in kwaZulu Natal", *Alternation*, 10, pp. 229–257, 2003.
- S.J. Cowley, "Contextualizing bodies: how human responsiveness constrains distributed cognition", *Language Sci.*, 26, pp. 565–591, 2004.
- S.J. Cowley and T. Kanda, "Friendly machines: interaction-oriented robots today and tomorrow", *Alternation*, 12, pp. 79–106, 2005.
- S.J. Cowley and K.F. MacDorman, "Simulating conversations: the communion game", *AI & Soc.*, 9, pp. 116–137, 1995.
- S.J. Cowley, S. Moodley and A. Fiori-Cowley, "Grounding signs of culture: primary intersubjectivity in social semiosis", *Mind, Culture Activity*, 11, pp. 109–132, 2004.
- R. Cummins, "Functional analysis", *J. Phil.*, 72, pp. 741–760, 1975.
- A. Damasio, *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*, London: Heinemann, 1999.
- P.S. Davies, *Norms of Nature: Naturalism and the Nature of Functions*, Cambridge MA: MIT Press, 2001.
- T. Deacon, *The Symbolic Species*, Harmondsworth: Penguin, 1997.
- D.C. Dennett, *Consciousness Explained*, Boston: Little, Brown & Company, 1991a.
- D.C. Dennett, "Real patterns", *J. Phil.*, 88, pp. 27–51, 1991b.
- P. Dourish, "What we talk about when we talk about context", *Pers. Ubiquitous Comput.*, 8, pp. 19–30, 2004.
- J.A. Fodor and Z.W. Pylyshyn, "Connectionism and cognitive architecture: a critical analysis", *Cognition*, 28, pp. 3–71, 1988.
- E. Goffman, "On face-work: an analysis of ritual elements in social interaction", *Psychiatry: J. Interpersonal Relat.*, 18, pp. 213–231, 1955.
- E. Goffman, *Behavior in Public Places: Notes on the Social Organization of Gatherings*, New York: Free Press, 1963.
- E. Goffman, "The interaction order", *Am. Sociol. Rev.*, 48, pp. 1–17, 1983.
- S. Harnad, "The symbol grounding problem", *Physica D*, 42, pp. 335–346, 1990a.
- S. Harnad, "Lost in the hermeneutic hall of mirrors", *J. Exp. Theor. Artif. Intell.*, 2, pp. 321–327, 1990b.
- M.D. Hauser, N. Chomsky and W.T. Fitch, "The faculty of language: What is it, who has it, and how did it evolve?", *Science*, 298, pp. 1569–1579, 2002.
- R.A. Hinde, *Individuals, Relationships and Culture: Links between Ethology and the Social Sciences*, Cambridge: Cambridge University Press, 1987.
- E. Hutchins, "How a cockpit remembers its speed", *Cognitive Sci.*, 19, pp. 265–288, 1995.
- H. Ishiguro, "Android science: conscious and subconscious recognition", *Connection Sci.*, 2006 (this issue).
- L.-E. Janlert, "The frame problem: Freedom or stability? With pictures we can have both", in *The Robot's Dilemma Revisited: The Frame Problem in Artificial Intelligence*, K.M. Ford and Z.W. Pylyshyn, Eds, Norwood, NJ: Ablex, pp. 35–48, 1996.
- P. Johnson-Laird, *The Computer and the Mind*, Cambridge, MA: Harvard University Press, 1988.
- K. Kaye, *The Mental and Social Life of Babies: How Parents Create Persons*, London: Methuen, 1982.
- D. Kirsh and P. Maglio, "Some epistemic benefits of action: Tetris a case study", in *Proceedings of the Fourteenth Annual Cognitive Science Society*, San Francisco: Morgan Kaufmann, 1992.
- D. Kirsh and P. Maglio, "On distinguishing epistemic from pragmatic action", *Cognitive Sci.*, 18, pp. 513–549, 1994.
- G. Kugiumutzakis, "Genesis and development of mimesis in early imitation of early facial and vocal models", in *Imitation in Infancy*, G. Nadel and G. Butterworth, Eds, Cambridge: Cambridge University Press, pp. 36–59, 1999.
- H. Kummer, *In Quest of the Sacred Baboon*, M.A. Biederman-Thorson, Transl., Princeton, NJ: Princeton University Press, 1995.
- B. Lee, "Empathy, androids and 'authentic experience'", *Connection Sci.*, 2006 (this issue).
- M. Legerstee, *Infants' Sense of People: Precursors to a Theory of Mind*, Cambridge: Cambridge University Press, 2005.
- J. Lindblom and T. Ziemke, "The social body in motion: cognitive development in infants and androids", *Connection Sci.*, 2006 (this issue).
- K.F. MacDorman, "Grounding symbols through sensorimotor integration", *J. Robot. Soc. Jpn.*, 17, pp. 20–24, 1999.
- K.F. MacDorman and S.J. Cowley, "Long-term relationships as a benchmark for robot personhood", in *Proceedings of the 15th IEEE International Symposium on Robot and Human Interactive Communication*, 2006.

- K.F. MacDorman and H. Ishiguro, "The study of interaction through the development of androids". Computer Vision and Image Processing Workshop, Information Processing Society of Japan, SIG Technical Reports 2004-CVIM-146, pp. 69–75 (2004).
- K.F. MacDorman and H. Ishiguro, "The uncanny advantage of using androids in social and cognitive science research", *Interact. Stud.*, 7(3), pp. 297–337, 2006a.
- K.F. MacDorman and H. Ishiguro, "Opening Pandora's uncanny box", *Interact. Stud.*, 7(3), pp. 361–368, 2006b.
- K.F. MacDorman, T. Minato, M. Shimada, S. Itakura, S.J. Cowley and H. Ishiguro, "Assessing human likeness by eye contact in an android testbed", in *Proceedings of the XXVII Annual Meeting of the Cognitive Science Society*, 2005.
- K.F. MacDorman, K. Tatani, Y. Miyazaki and M. Koeda, "Proto-symbol emergence", in *Proceedings of The IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2000.
- B.B. Mangan, "Sensation's ghost: the non-sensory 'fringe' of consciousness", *Psyche*, 7(18), 2001.
- E. Marais, *My Friends the Baboons*, London: Bond & Briggs, 1939.
- J. McCarthy and P.J. Hayes, "Some philosophical problems from the standpoint of artificial intelligence", in *Machine Intelligence 4*, B. Meltzer and D. Michie, Eds, Edinburgh: Edinburgh University Press, pp. 463–502, 1969.
- A. Melzoff and M. Moore, "Explaining facial imitation: a theoretical model", *Early Dev. Parent.*, 6, pp. 179–192, 1997.
- A. Newell, "Physical symbol systems", *Cognitive Sci.*, 4, pp. 135–183, 1980.
- Z. Pylyshyn, *Computation and Cognition*, Cambridge, MA: MIT Press, 1984.
- P. Rochat, *The Infant's World*, Cambridge, MA: Harvard University Press, 2000.
- D. Ross and P. Dumouchel, "Emotions as strategic signals", *Rational. Soc.*, 16, pp. 251–286, 2004.
- T. Sawada, T. Takagi and M. Fujita, "Behavior selection and motion modulation in emotionally grounded architecture for QRIO SDR-4X II", in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2514–2519, 2004.
- J. Searle, "Minds, brains, and programs", *Behav. Brain Sci.*, 3, pp. 417–424, 1980.
- D. Stern, *The First Relationship*, London: Fontana, 1977.
- T. Taylor, *Theorizing Language*, Oxford: Pergamon Press, 1997.
- G. Therborn, "Back to norms! On the scope and dynamics of norms and normative action", *Curr. Sociol.*, 50, pp. 863–880, 2002.
- M. Tomasello, "Beyond formalities: the case of language acquisition", *Linguist. Rev.*, 22, pp. 167–181, 2005.
- C. Trevarthen, "Communication and co-operation in early infancy: a description of primary intersubjectivity", in *Before Speech*, M. Bullowa, Ed., Cambridge: Cambridge University Press, pp. 321–347, 1979.
- C. Trevarthen, "Universal co-operative motives: how infants begin to know the language and skills of the language of their parents", in *Acquiring Culture*, G. Jahoda and I. Lewis, Eds, Beckenham: Croom Helm, pp. 37–90, 1988.
- C. Trevarthen, "The concept and foundations of infant intersubjectivity", in *Intersubjective Communication and Emotion in Early Ontogeny*, S. Bråten, Ed., Cambridge: Cambridge University Press, pp. 15–46, 1998.
- A.M. Turing, "Computing machinery and intelligence", *Mind*, 49, pp. 433–460, 1950.
- S. Turkle, W. Taggart, C.D. Kidd and O. Dasté, "Relational artifacts with children and elders: the complexities of cybercompanionship", *Connection Sci.*, 2006 (this issue).
- J. Weizenbaum, "Eliza: a computer program for the study of natural language communication between man and machine", *Commun. ACM*, 9, pp. 36–45, 1966.
- M. Wheeler, *Reconstructing the Cognitive World*, Cambridge, MA: MIT Press, 2005.
- L. Wittgenstein, *Phil. Investigat.*, Oxford: Blackwell, 1953.
- G.H. von Wright, *Norm and Action*, London: Routledge, 1963.
- L. Wright, "Functions", *Phil. Rev.*, 82, pp. 139–168, 1973.
- L. Wright, *Teleological Explanations*, Berkeley, CA: University of California Press, 1976.
- T. Ziemke and J. Lindblom, "Some methodological issues in android science", *Interact. Stud.*, 7(3), 2006.