Alternative Essences of Intelligence

Rodney A. Brooks, Cynthia Breazeal (Ferrell), Robert Irie, Charles C. Kemp, Matthew Marjanović, Brian Scassellati, Matthew M. Williamson
MIT Artificial Intelligence Lab, Cambridge, MA, 02139, USA
{brooks,ferrell,irie,cckemp,maddog,scaz,matt}@ai.mit.edu

Abstract

We present a novel methodology for building human-like artificially intelligent systems. We take as a model the only existing systems which are universally accepted as intelligent: humans. We emphasize building intelligent systems which are not masters of a single domain, but, like humans, are adept at performing a variety of complex tasks in the real world. Using evidence from cognitive science and neuroscience, we suggest four alternative essences of intelligence to those held by classical AI. These are the parallel themes of development, social interaction, embodiment, and integration. Following a methodology based on these themes, we have built a physical humanoid robot. In this paper we present our methodology and the insights it affords for facilitating learning, simplifying the computation underlying rich behavior, and building systems that can scale to more complex tasks in more challenging environments.

Introduction

An early development in the history of AI was the claim of Newell & Simon (1961) that humans use physical symbol systems to “think”. Over time, this has become adopted into Artificial Intelligence as an implicit and dominant hypothesis (see (Brooks 1991*a*) for a review). Although this assumption has begun to soften in recent years, a typical AI system still relies on uniform, explicit, internal representations of capabilities of the system, the state of the outside world, and the desired goals. These systems are dominated by search problems to both access the relevant facts, and determine how to apply them. Neo-classical AI adds Bayesian or other probabilistic ideas to this basic framework (Pearl 1988).

The underlying assumption of these approaches is that because a description of reasoning/behavior/-learning is possible at some level, then that description must be made explicit and internal to any mechanism that carries out the reasoning/behavior/learning. The realization that descriptions and mechanisms could be separated was one of the great breakthroughs of Rosenschein & Kaelbling (1986), but unfortunately that realization has been largely ignored. This introspective confusion between surface observations and deep structure has led AI away from its original goals of building complex, versatile, intelligent systems and towards the construction of systems capable of performing only within limited problem domains and in extremely constrained environmental conditions.

In this paper we present a methodology based on a different set of basis assumptions. We believe that human intelligence is a direct result of four intertwined attributes: developmental organization, social interaction, embodiment and physical coupling, and multimodal integration. Development forms the framework by which humans successfully acquire increasingly more complex skills and competencies. Social interaction allows humans to exploit other humans for assistance, teaching, and knowledge. Embodiment and physical coupling allow humans to use the world itself as a tool for organizing and manipulating knowledge. Integration allows humans to maximize the efficacy and accuracy of complementary sensory and motor systems.

We have followed this methodology to construct physical humanoid robots (see Figure 1). We design these robots to follow the same sorts of developmental paths that humans follow, building new skills upon earlier competencies. People interact with these robots through behavioral coupling and direct physical contact. The variety of sensory and motor systems on the robots provide ample opportunity to confront integration issues.

Using evidence from human behavior, and early results from our own work, we argue that building systems in this manner affords key insights into how to simplify the computation underlying rich behavior, how to facilitate learning, and how to create mechanisms that can scale to more complex tasks in more challenging environments.

The next section of this paper explores the assumptions about human intelligence which are deeply embedded within classical AI. The following sections explain how our methodology yields a plausible approach to creating robustly functioning intelligent systems, draw-

Figure 1: Our humanoid robot has undergone many transformations over the last few years. This is how it currently appears.

ing on examples from our research. The final section presents an outline of the key challenges to be faced along this new road in AI.

Assumptions about Intelligence

In recent years, AI research has begun to move away from the assumptions of classical AI: monolithic internal models, monolithic control, and general purpose processing. However, these concepts are still prevalent in much current work, and are deeply ingrained in many architectures for intelligent systems. For example, in the recent AAAI-97 Proceedings, one sees a continuing interest in planning ((Littman 1997, Hauskrecht 1997, Boutilier & Brafman 1997, Blythe & Veloso 1997, Brafman 1997)) and representation ((McCain & Turner 1997, Costello 1997, Lobo, Mendez & Taylor 1997)), which build on these assumptions.

The motivation for our alternative methodology comes from a modern understanding of cognitive science and neuroscience, which counterposes the assumptions of classical AI, as described in the following sections.

Humans have no full monolithic internal models. There is evidence that in normal tasks humans tend to minimize their internal representation of the world. Ballard, Hayhoe & Pelz (1995) have shown that in performing a complex task, like building a copy of a display of blocks, humans do not build an internal model of the entire visible scene. By changing the display while subjects were looking away, Ballard found that subjects noticed only the most drastic of changes; rather than keeping a complete model of the scene, they instead left that information in the world and continued to refer back to the scene while performing the copying task.

There is also evidence that there are multiple internal representations, which are not mutually consistent. For example, in the phenomena of blindsight, cortically blind patients can discriminate different visual stimuli, but actually report seeing nothing (Weiskrantz 1986). This inconsistency would not be a feature of a single central model of visual space.

These experiments and many others like it (e.g. (Rensink, O'Regan & Clark 1997, Gazzaniga & LeDoux 1978)) convincingly demonstrate that humans do not construct a full, monolithic model of the environment. Instead humans tend to only represent what is immediately relevant from the environment, and those representations do not have full access to one another.

Humans have no monolithic control. Naive introspection and observation can lead one to believe in a neurological equivalent of the central processing unit – something that makes the decisions and controls the other functions of the organism. While there are undoubtedly control structures, this model of a single, unitary control system is not supported by evidence from cognitive science.

One example comes from studies of split brain patients by Gazzaniga & LeDoux (1978). These are patients where the corpus callosum (the main structure connecting the two hemispheres of the brain) has been cut. The patients are surprisingly normal after the operation, but with deficits that are revealed by presenting different information to either side of the (now unconnected) brain. Since each hemisphere controls one side of the body, the experimenters can probe the behavior of each hemisphere independently (for example, by observing the subject picking up an object appropriate to the scene that they had viewed). In one example, a snow scene was presented to the right hemisphere and the leg of a chicken to the left. The subject selected a chicken head to match the chicken leg, explaining with the verbally dominant left hemisphere that "I saw the claw and picked the chicken". When the right hemisphere then picked a shovel to correctly match the snow, the left hemisphere explained that you need a shovel to "clean out the chicken shed" (Gazzaniga & LeDoux 1978, p.148). The separate halves of the subject independently acted appropriately, but one side falsely explained the choice of the other. This suggests that there are multiple independent control systems, rather than a single monolithic one.

Humans are not general purpose. The brain is conventionally thought to be a general purpose machine, acting with equal skill on any type of operation that it performs by invoking a set of powerful rules. However, humans seem to be proficient only in particular sets of skills, at the expense of other skills, often in non-obvious ways. A good example of this is the Stroop effect (Stroop 1935). When presented with a list of words written in a variety of colors, performance in a color recognition and articulation task is actually dependent on the semantic content of the words; the task is very difficult if names of colors are printed in non-corresponding colors. This experiment demonstrates the specialized nature of human computational processes and interactions.

Even in the areas of deductive logic, humans often perform extremely poorly in different contexts. Wason (1966) found that subjects were unable to apply the negative rule of if-then inference when four cards were labeled with single letters and digits. However, with additional context—labeling the cards such that they were understandable as names and ages—subjects could easily solve exactly the same problem.

Further, humans often do not use subroutine-like rules for making decisions. They are often more emotional than rational, and there is evidence that this emotional content is an important aspect of decision making (Damasio 1994).

Essences of Human Intelligence

Since humans are vastly complex systems, we do not expect to duplicate every facet of their operation. However, we must be very careful not to ignore aspects of human intelligence solely because they appear complex. Classical and neo-classical AI tends to ignore or avoid these complexities, in an attempt to simplify the problem (Minsky & Papert 1970). We believe that many of these discarded elements are essential to human intelligence and that they actually simplify the problem of creating human-like intelligence.

Development Humans are not born with complete reasoning systems, complete motor systems, or even complete sensory systems. Instead, they undergo a process of development where they are able to perform more difficult tasks in more complex environments en route to the adult state. This is a gradual process, in which earlier forms of behavior disappear or are modified into more complex types of behavior. The adaptive advantage of the earlier forms appears to be that they prepare and enable more advanced forms of behavior to develop within the situated context they provide. The developmental psychology literature abounds with examples of this phenomenon. For instance, the work of Diamond (1990) shows that infants between five and twelve months of age progress through a number of distinct phases in the development of visually guided reaching. In one reaching task, the infant must retrieve a toy from inside a transparent box with only one open side. In this progression, infants in later phases consistently demonstrate more sophisticated reaching strategies to retrieve the toy in more challenging scenarios. As the infant's reaching competency develops, later stages incrementally improve upon the competency afforded by the previous stage.

Social Interaction Human infants are extremely dependent on their caregivers, relying upon them not only for basic necessities but also as a guide to their development. The presence of a caregiver to nurture the child as it grows is essential. This reliance on social contact is so integrated into our species that it is hard to imagine a completely asocial human. However, severe developmental disorders sometimes give us a glimpse of the importance of social contact. One example is autism. Autistic children often appear completely normal on first examination; they look normal, have good motor control, and seem to have normal perceptual abilities. However, their behavior is completely strange to us, in part because they do not recognize or respond to normal social cues (Baron-Cohen 1995). They do not maintain eye contact, recognize pointing gestures, or understand simple social conventions. Even the most highly functioning autistics are severely disabled in our society.

Embodiment Perhaps the most obvious, and most overlooked, aspect of human intelligence is that it is embodied. Humans are embedded in a complex, noisy, constantly changing environment. There is a direct physical coupling between action and perception, without the need for an intermediary representation. This coupling makes some tasks simple and other tasks more complex. By exploiting the properties of the complete system, certain seemingly complex tasks can be made computationally simple. For example, when putting a jug of milk in the refrigerator, you can exploit the pendulum action of your arm to move the milk (Greene 1982). The swing of the jug does not need to be explicitly planned or controlled, since it is the natural behavior of the system. Instead of having to plan the whole motion, the system only has to modulate, guide and correct the natural dynamics. For an embodied system, internal representations can be ultimately grounded in sensory-motor interactions with the world (Lakoff 1987).

Integration Humans have the capability to receive an enormous amount of information from the world. Visual, auditory, somatosensory, and olfactory cues are all processed simultaneously to provide us with our view of the world. However, there is evidence that the sensory modalities are not independent; stimuli from one modality can and do influence the perception of stimuli in another modality. Churchland, Ramachandran & Sejnowski (1994) describe an experiment illustrating how audition can cause illusory visual motion. A fixed square and a dot located to its left are presented to the observer. Without any sound stimuli, the blinking of the dot does not result in any perception of motion. If a tone is alternately played in the left and right ears, with the left ear tone coinciding with the dot pre-

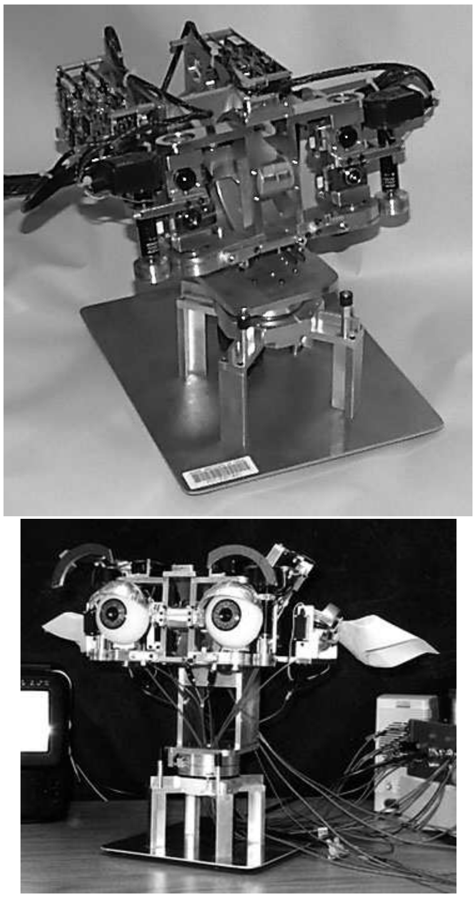

Figure 2: We have built two active vision heads, similar in design to Cog's head. On top is a desktop version with a 1 DOF neck, and below a head with actuators to include facial expressions.

sentation, there is an illusory perception of back and forth motion of the dot, with the square acting as a visual occluder. Vision can cause auditory illusions too, such as the McGurk effect (Cohen & Massaro 1990). These studies demonstrate that humans' perception of their senses cannot be treated as completely independent processes.

Methodology

Our methodology—exploring themes of development, social interaction, physical interaction and integration while building real robots—is motivated by two ideas. First, we believe that these themes are important aspects of human intelligence. Second, from an engineering perspective, these themes make the problems of building human intelligence easier.

Embodiment

A principle tenet of our methodology is to build and test real robotic systems. We believe that building human-like intelligence requires human-like interaction with the world (Brooks & Stein 1994). Humanoid form is important to allow humans to interact with the robot in a natural way. In addition we believe that building a real system is computationally less complex than simulating such a system. The effects of gravity, friction, and natural human interaction are obtained for free, without any computation.

Our humanoid robot, named Cog and shown in Figure 1, approximates a human being from the waist up with twenty-one degrees-of-freedom (DOF) and a variety of sensory systems. The physical structure of the robot, with movable torso, arms, neck and eyes gives it human-like motion, while the sensory systems (visual, auditory, vestibular, and proprioceptive) provide rich information about the robot and its immediate environment. These together present many opportunities for interaction between the robot and humans.

In addition to the full humanoid, we have also developed active head platforms, of similar design to Cog's head, as shown in Figure 2 (Scassellati 1998*a*). These self-contained systems allow us to concentrate on various issues in close human-machine interaction, including face detection, imitation, emotional display and communication, etc. (Scassellati 1998*b*, Ferrell 1998*c*).

Development

Building systems developmentally facilitates learning both by providing a structured decomposition of skills and by gradually increasing the complexity of the task to match the competency of the system.

Bootstrapping Development is an *incremental* process. As it proceeds, prior structures and their behavioral manifestations place important constraints on the later structures and proficiencies. The earlier forms bootstrap the later structures by providing subskills and knowledge that can be re-used. By following the developmental progression, the learning difficulties at each stage are minimized. Within our group, Scassellati (1996) discusses how a humanoid robot might acquire basic social competencies through this sort of developmental methodology.

The work of Marjanović, Scassellati & Williamson (1996) applied bootstrapping techniques to our robot, coordinating visual and motor systems by learning to point toward a visual target. A map used for a saccading behavior (visual/eye-movement map), was reused to learn a reaching behavior (visual/arm-movement map). The learned saccadic behavior bootstrapped the reaching behavior, reducing the complexity of the overall learning task. Other examples of developmental learning that we have explored can be found in (Ferrell 1996).

Gradual increase in complexity The developmental process, starting with a simple system that gradually becomes more complex allows efficient learning throughout the whole process. For example, infants are born with low acuity vision which simplifies the visual input they must process. The infant's visual performance develops in step with their ability to process the influx of stimulation (Johnson 1993). The same is true for the motor system. Newborn infants do not have independent control over each degree of freedom of their limbs, but through a gradual increase in the granularity of their motor control they learn to coordinate the full complexity of their bodies. A process where the acuity of both sensory and motor systems are gradually increased significantly reduces the difficulty of the learning problem (Thelen & Smith 1994).

To further facilitate learning, the gradual increase in internal complexity associated with development should be accompanied by a gradual increase in the complexity of the external world. For an infant, the caregiver biases how learning proceeds by carefully structuring and controlling the complexity of the environment. This approach is in stark contrast to most machine learning methods, where the robot learns in a usually hostile environment, and the bias, instead of coming from the robots' interaction with the world, is included by the designer. We believe that gradually increasing the complexity of the environment makes learning easier and more robust.

By exploiting a gradual increase in complexity both internal and external, while reusing structures and information gained from previously learned behaviors, we hope to be able to learn increasingly sophisticated behaviors. We believe that these methods will allow us to construct systems which do scale autonomously (Ferrell & Kemp 1996).

Social Interaction

Building social skills into an artificial intelligence provides not only a natural means of human-machine interaction but also a mechanism for bootstrapping more complex behavior. Our research program has investigated social interaction both as a means for bootstrapping and as an instance of developmental progression.

Bootstrapping Social interaction can be a means to facilitate learning. New skills may be socially transfered from caregiver to infant through mimicry or imitation, through direct tutelage, or by means of scaffolding, in which a more able adult manipulates the infant's interactions with the environment to foster novel abilities. Commonly scaffolding involves reducing distractions, marking the task's critical attributes, reducing the number of degrees of freedom in the target task, and enabling the subject to experience the end or outcome before the infant is cognitively or physically able of seeking and attaining it for herself (Wood, Bruner & Ross 1976).

We are currently engaged in work studying bootstrapping new behaviors from social interactions. One research project focuses on building a robotic system capable of learning communication behaviors in a social context where the human provides various forms of scaffolding to facilitate the robot's learning task (Ferrell 1998*b*). The system uses expressive facial gestures (see Figure 2) as feedback to the caregiver (Ferrell 1998*a*). The caregiver can then regulate the complexity of the social interaction to optimize the robot's learning rate.

Development of social interaction The social skills required to make use of scaffolding are complex. Infants acquire these social skills through a developmental progression (Hobson 1993). One of the earliest precursors is the ability to share attention with the caregiver. This ability can take many forms, from the recognition of a pointing gesture to maintaining eye contact.

In our work, we have also examined social interaction from this developmental perspective. One research program focuses on a developmental implementation of shared attention mechanisms based upon normal child development, developmental models of autism[1], and on models of the evolutionary development of social skills (Scassellati 1996). The first step in this developmental progression is recognition of eye contact. Human infants are predisposed to attend to socially relevant stimuli, such as faces and objects that have human-like motion. The system is currently capable of detecting faces in its peripheral vision, saccading to the faces, and finding eyes within its foveal vision (Scassellati 1998*b*). This developmental chain has also produced a simple imitation behavior; the head will mimic yes/no head nods of the caregiver (Scassellati 1998*c*).

Physical Coupling

Another aspect of our methodology is to exploit interaction and tight coupling between the robot and its environment to give complex behavior, to facilitate learning, and to avoid the use of explicit models. Our systems are physically coupled with the world and operate directly in that world without any explicit representations of it (Brooks 1986, Brooks 1991*b*). There are representations, or accumulations of state, but these only refer to the internal workings of the system; they are meaningless without interaction with the outside world. The embedding of the system within the world enables the internal accumulations of state to provide useful behavior (this was the fundamental approach taken by Ashby (1960) contemporaneously with the development of early AI).

One example of such a scheme is implemented to control our robot's arms. As detailed in (Williamson 1998*a*, Williamson 1998*b*), a set of self-adaptive oscillators are used to drive the joints of the arm. Each joint is actuated by a single oscillator, using proprioceptive information at that joint to alter the frequency and phase of the joint motion. There are no connections between the oscillators, except indirectly through

[1] Some researchers believe that the missing mechanisms of shared attention may be central to autism disorders (Baron-Cohen 1995). In comparison to other mental retardation and developmental disorders (like Williams and Downs Syndromes), the deficiencies of autism in this area are quite specific (Karmiloff-Smith, Klima, Bellugi, Grant & Baron-Cohen 1995).

the physical structure of the arm. Without using any kinematic or dynamical models, this simple scheme has been used for a variety of different coordinated tasks, including turning a crank and playing with a slinky toy. The interaction between the arm and the environment enables the oscillators to generate useful behavior. For example, without the slinky to connect the two arms, they are uncoordinated, but when the arms are coupled through the slinky, the oscillators tune into the dynamics of the motion and coordinated behavior is achieved. In all cases, there is no central controller, and no modeling of the arms or the environment; the behavior of the whole system comes from the coupling of the arm and controller dynamics. Other researchers have built similar systems which exhibit complex behavior with either simple or no control (McGeer 1990, Schaal & Atkeson 1993) by exploiting the system dynamics.

Sensory Integration

Sensory Integration Simplifies Computation Some tasks are best suited for particular sensory modalities. Attempting to perform the task using only one modality is sometimes awkward and computationally intensive. Utilizing the complementary nature of separate modalities results in a reduction of overall computation. We have implemented several mechanisms on Cog that use multimodal integration to aid in increasing performance or developing competencies.

For example, Peskin & Scassellati (1997) implemented a system that stabilized images from a moving camera using vestibular feedback. Rather than attempting to model the camera motion, or to predict motion effects based on efference copy, the system mimics the human vestibular-ocular reflex (VOR) by compensating for camera motion through learned feedback from a set of rate gyroscopes. By integrating two sensory systems, we can achieve better performance than traditional image processing methods, while using less computation.

Sensory Integration Facilitates Learning By integrating different sensory modalities we can exploit the complex nature of stimuli to facilitate learning. For example, objects that make noise often move. This correlation can be exploited to facilitate perception. In our work, we have investigated primarily the relationship between vision and audition in learning to orient toward stimuli.

We can characterize this relationship by examining developmental evidence. Wertheimer (1961) has shown that vision and audition interact from birth; even ten-minute-old children will turn their eyes toward an auditory cue. This interaction between the senses continues to develop, indeed related investigations with young owls have determined that visual stimuli greatly affect the development of sound localization. With a constant visual bias from prisms, owls adjusted their sound localization to match the induced visual errors (Knudsen & Knudsen 1985).

Irie (1997) built an auditory system for our robot that utilizes visual information to train auditory localization; the visually-determined location of a sound source with a corresponding motion is used to train an auditory spatial map. This map is then used to orient the head toward the object. This work highlights not only the development of sensory integration, but also the simplification of computational requirements that can be obtained through integration.

Challenges for Intelligent Systems

This new approach to designing and studying intelligent systems leaves us with a new set of challenges to overcome. Here are the key questions which we must now answer.

- Scaling and Development: What learning structures and organizational principles will allow us to design successful developmental systems?
- Social Interaction:
  - How can the system learn to communicate with humans?
  - What attributes of an infant must the machine emulate in order to elicit caregiver behavior from humans?
  - What drives, motivations and emotions are necessary for the system to communicate effectively with humans?
- Physical Coupling:
  - How can a system scale the complexity of its coordinated motion while still exploiting the dynamics?
  - How can the system combine newly learned spatial skills with previously learned spatial skills? How is that memory to be organized?
  - How can the system use previously learned skills in new contexts and configurations?
- Integration:
  - How can a conglomerate of subsystems, each with different or conflicting goals and behaviors, act with coherence and stability?
  - At what scale should we emulate the biological organism, keeping in mind engineering constraints?
  - What are good measures of performance for integrated systems which develop, interact socially and are physically coupled with the world?

Conclusion

We have reported on a work in progress which incorporates a new methodology for achieving Artificial Intelligence. We have built a humanoid robot that operates and develops in the world in ways that are similar to the ways in which human infants operate and develop. We have demonstrated learning to saccade, learning to correlate auditory and occular coordinate systems, self adapting vestibular-ocular systems, coordinated neck and ocular systems, learning of hand-eye coordination, localization of multiple sounds streams, variable stiffness arms that interact safely with people, arm control based on biological models of invertebrate spinal circuits, adaptive arm control that tunes into to subtle physical cues from the environment, face detection, eye detection to find gaze direction, coupled human robot interaction that is a precursor to caregiver scaffolding for learning, large scale touch sensitive body skin, and multi-fingered hands that learn to grasp objects based on self-categorized stiffness properties. These are components for higher level behaviors that we are beginning to put together using models of shared attention, emotional coupling between robot and caregiver, and developmental models of human infants.

We have thus chosen to approach AI from a different perspective, in the questions we ask, the problems we try to solve, and the methodology and inspiration we use to achieve our goals. Our approach has led us not only to formulate the problem in the context of general human-level intelligence, but to redefine the essences of that intelligence. Traditional AI work has sought to narrowly define a problem and apply abstract rules to its solution. We claim that our goal of creating a learning, scalable, intelligent system, with competencies similar to human beings, is altogether relevant in trying to solve a broad class of real-world situated problems. We further believe that the principles of development, social interaction, physical coupling to the environment, and integration are essential to guide us towards our goal.

Acknowledgments

This project is partially funded by an ONR/ARPA Vision MURI Grant (No. N00014-95-1-0600).

References

Ashby, W. R. (1960), *Design for a Brain*, second edn, Chapman and Hall.

Ballard, D., Hayhoe, M. & Pelz, J. (1995), 'Memory representations in natural tasks', *Journal of Cognitive Neuroscience* pp. 66–80.

Baron-Cohen, S. (1995), *Mindblindness*, MIT Press.

Blythe, J. & Veloso, M. (1997), Analogical Replay for Efficient Conditional Planning, *in* 'Proceedings of American Association of Artificial Intelligence (AAAI-97)', pp. 668–673.

Boutilier, C. & Brafman, R. I. (1997), Planning with Concurrent Interacting Actions, *in* 'Proceedings of American Association of Artificial Intelligence (AAAI-97)', pp. 720–726.

Brafman, R. I. (1997), A Heuristic Variable Grid Solution Method for POMDPs, *in* 'Proceedings of American Association of Artificial Intelligence (AAAI-97)', pp. 727–733.

Brooks, R. A. (1986), 'A Robust Layered Control System for a Mobile Robot', *IEEE Journal of Robotics and Automation* **RA-2**, 14–23.

Brooks, R. A. (1991*a*), Intelligence Without Reason, *in* 'Proceedings of the 1991 International Joint Conference on Artificial Intelligence', pp. 569–595.

Brooks, R. A. (1991*b*), 'Intelligence Without Representation', *Artificial Intelligence Journal* **47**, 139–160. originally appeared as MIT AI Memo 899 in May 1986.

Brooks, R. A. & Stein, L. A. (1994), 'Building brains for bodies', *Autonomous Robots* **1**(1), 7–25.

Churchland, P., Ramachandran, V. & Sejnowski, T. (1994), A Critique of Pure Vision, *in* C. Koch & J. Davis, eds, 'Large-Scale Neuronal Theories of the Brain', MIT Press.

Cohen, M. & Massaro, D. (1990), 'Synthesis of visible speech', *Behaviour Research Methods, Intruments and Computers* **22**(2), pp. 260–263.

Costello, T. (1997), Beyond Minimizing Change, *in* 'Proceedings of American Association of Artificial Intelligence (AAAI-97)', pp. 448–453.

Damasio, A. R. (1994), *Descartes' Error*, G.P. Putnam's Sons.

Diamond, A. (1990), Developmental Time Course in Human Infants and Infant Monkeys, and the Neural Bases of Inhibitory Control in Reaching, *in* 'The Development and Neural Bases of Higher Cognitive Functions', Vol. 608, New York Academy of Sciences, pp. 637–676.

Feigenbaum, E. A. & Feldman, J., eds (1963), *Computers and Thought*, McGraw-Hill, New York.

Ferrell, C. (1996), Orientation Behavior using Registered Topographic Maps, *in* 'From Animals to Animats: Proc 1996 Society of Adaptive Behavior', Cape Cod, Massachusetts, pp. 94–103.

Ferrell, C. (1998*a*), Emotional Robots and Learning During Social Exchanges. Submitted to Autonomous Agents-98.

Ferrell, C. (1998*b*), 'Learning by Scaffolding'. MIT Ph.D. Thesis Proposal.

Ferrell, C. (1998*c*), A Motivational System for Regulating Human-Robot Interaction. Submitted to AAAI-98.

Ferrell, C. & Kemp, C. (1996), An Ontogenetic Perspective to Scaling Sensorimotor Intelligence, *in* 'Embodied Cognition and Action: Papers from the 1996 AAAI Fall Symposium', AAAI Press.

Gazzaniga, M. S. & LeDoux, J. E. (1978), *The Integrated Mind*, Plenum Press, New York.

Greene, P. H. (1982), 'Why is it easy to control your arms?', *Journal of Motor Behavior* **14**(4), 260–286.

Hauskrecht, M. (1997), Incremental Methods for computing bounds in partially observable Markov decision processes, *in* 'Proceedings of American Association of Artificial Intelligence (AAAI-97)', pp. 734–739.

Hobson, R. P. (1993), *Autism and the Development of Mind*, Erlbaum.

Irie, R. E. (1997), Multimodal Sensory Integration for Localization in a Humanoid Robot, *in* 'Proceedings of Second IJCAI Workshop on Computational Auditory Scene Analysis (CASA'97)', IJCAI-97.

Johnson, M. H. (1993), Constraints on Cortical Plasticity, *in* M. H. Johnson, ed., 'Brain Development and Cognition: A Reader', Blackwell, Oxford, pp. 703–721.

Karmiloff-Smith, A., Klima, E., Bellugi, U., Grant, J. & Baron-Cohen, S. (1995), 'Is there a social module? Language, face processing, and theory of mind in individuals with Williams Syndrome', *Journal of Cognitive Neuroscience* **7:2**, 196–208.

Knudsen, E. I. & Knudsen, P. F. (1985), 'Vision Guides the Adjustment of Auditory Localization in Young Barn Owls', *Science* **230**, 545–548.

Lakoff, G. (1987), *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*, University of Chicago Press, Chicago, Illinois.

Littman, M. L. (1997), Probabilistic Propositional Planning: Representations and Complexity, *in* 'Proceedings of American Association of Artificial Intelligence (AAAI-97)', pp. 748–754.

Lobo, J., Mendez, G. & Taylor, S. R. (1997), Adding Knowledge to the Action Description Language $\mathcal{A}$, *in* 'Proceedings of American Association of Artificial Intelligence (AAAI-97)', pp. 454–459.

Marjanović, M. J., Scassellati, B. & Williamson, M. M. (1996), Self-Taught Visually-Guided Pointing for a Humanoid Robot, *in* 'From Animals to Animats: Proceedings of 1996 Society of Adaptive Behavior', Cape Cod, Massachusetts, pp. 35–44.

McCain, N. & Turner, H. (1997), Causal Theories of Action and Change, *in* 'Proceedings of American Association of Artificial Intelligence (AAAI-97)', pp. 460–465.

McGeer, T. (1990), Passive Walking with Knees, *in* 'Proc 1990 IEEE Intl Conf on Robotics and Automation'.

Minsky, M. & Papert, S. (1970), 'Draft of a proposal to ARPA for research on artificial intelligence at MIT, 1970-71'.

Newell, A. & Simon, H. (1961), GPS, a program that simulates thought, *in* H. Billing, ed., 'Lernende Automaten', R. Oldenbourg, Munich, Germany, pp. 109–124. Reprinted in (Feigenbaum and Feldman, 1963, pp.279–293).

Pearl, J. (1988), *Probabilistic Reasoning in Intelligent Systems*, Morgan Kaufmann, San Mateo, CA.

Peskin, J. & Scassellati, B. (1997), Image Stabilization through Vestibular and Retinal Feedback, *in* R. Brooks, ed., 'Research Abstracts', MIT Artificial Intelligence Laboratory.

Rensink, R., O'Regan, J. & Clark, J. (1997), 'To See or Not to See: The Need for Attention to Perceive Changes in Scenes', *Psychological Science* **8**, 368–373.

Rosenschein, S. J. & Kaelbling, L. P. (1986), The Synthesis of Machines with Provable Epistemic Properties, *in* J. Halpern, ed., 'Proc. Conf. on Theoretical Aspects of Reasoning about Knowledge', Morgan Kaufmann Publishers, Los Altos, California, pp. 83–98.

Scassellati, B. (1996), Mechanisms of Shared Attention for a Humanoid Robot, *in* 'Embodied Cognition and Action: Papers from the 1996 AAAI Fall Symposium', AAAI Press.

Scassellati, B. (1998*a*), A Binocular, Foveated Active Vision System, Technical report, MIT Artificial Intelligence Lab. In submission.

Scassellati, B. (1998*b*), Finding Eyes and Faces with a Foveated Vision System, *in* 'Proceedings of the American Association of Artificial Intelligence'. Submitted to AAAI-98.

Scassellati, B. (1998*c*), Imitation and Mechanisms of Shared Attention: A Developmental Structure for Building Social Skills, *in* 'Agents in Interaction - Acquiring Competence through Imitation: Papers from a Workshop at the Second International Conference on Autonomous Agents'.

Schaal, S. & Atkeson, C. G. (1993), Open loop Stable Control Strategies for Robot Juggling, *in* 'Proceedings 1993 IEEE International Conference on Robotics and Automation', Vol. 3, pp. 913–918.

Stroop, J. (1935), 'Studies of interference in serial verbal reactions', *Journal of Experimental Psychology* **18**, 643–62.

Thelen, E. & Smith, L. (1994), *A Dynamic Systems Approach to the Development of Cognition and Action*, MIT Press, Cambridge, MA.

Wason, P. C. (1966), *New Horizons in Psychology*, Vol. 1, Penguin Books, Harmondsworth, England, pp. 135–51.

Weiskrantz, L. (1986), *Blindsight: A Case Study and Implications*, Clarendon Press, Oxford.

Wertheimer, M. (1961), 'Psychomotor coordination of auditory and visual space at birth', *Science* **134**, 1692.

Williamson, M. M. (1998*a*), Exploiting natural dynamics in robot control, *in* 'Fourteenth European Meeting on Cybernetics and Systems Research (EMCSR '98)', Vienna, Austria.

Williamson, M. M. (1998*b*), Rhythmic robot control using oscillators. Submitted to IROS '98.

Wood, D., Bruner, J. S. & Ross, G. (1976), 'The role of tutoring in problem-solving', *Journal of Child Psychology and Psychiatry* **17**, 89–100.