

Affective learning — a manifesto

R W Picard, S Papert, W Bender, B Blumberg, C Breazeal, D Cavallo, T Machover, M Resnick, D Roy and C Strohecker

The use of the computer as a model, metaphor, and modelling tool has tended to privilege the 'cognitive' over the 'affective' by engendering theories in which thinking and learning are viewed as information processing and affect is ignored or marginalised. In the last decade there has been an accelerated flow of findings in multiple disciplines supporting a view of affect as complexly intertwined with cognition in guiding rational behaviour, memory retrieval, decision-making, creativity, and more. It is time to redress the imbalance by developing theories and technologies in which affect and cognition are appropriately integrated with one another. This paper describes work in that direction at the MIT Media Lab and projects a large perspective of new research in which computer technology is used to redress the imbalance that was caused (or, at least, accentuated) by the computer itself.

1. Vision

The last half-century of technological acceleration has yielded a massive incursion of digital technology into the learning environment, making dramatic differences, and promising even greater changes, to the practice of learning. Computers have served as tools to aid in learning at all levels from simple classroom activities to the way theorists think about thinking. The field of artificial intelligence, with emphasis on ideas such as knowledge representation, modelling of logical processes, and other kinds of important cognitive activities, has prompted thinking about parallel concepts in human learning, and facilitated the development of theories where thinking and learning are viewed as information processing. Both human and machine learning research have benefited from this exchange of ideas between psychology and computation.

However, these benefits have been bought at the price of a bias towards the cognitive and relative neglect of the affective. Of course nobody denies the role of affect in learning. Certainly teachers know that it plays a crucial role, recognising it under intuitively understood headings like motivation, emotion, interest, and attention. Even leading theorists of the cognitive scientific revolution [1, 2] have called for greater representation of affect. However, the extension of cognitive theory to explain and exploit the role of affect in learning is at best in its infancy.

Developments in the past decade have both accentuated the gap in theoretical understanding between these two sides of mental functioning and offered glimpses of insights into how to close it. On the most fundamental level, an accelerated flow of findings in neuroscience, psychology, and cognitive science

itself present affect as complexly intertwined with thinking, and performing important functions with respect to guiding rational behaviour, memory retrieval, decision-making, creativity, and more. While it has always been understood that too much emotion is bad for rational thinking, recent findings suggest that so too is too little emotion — when basic mechanisms of emotion are missing in the brain, then intelligent functioning is hindered. These findings point to new advances in understanding the human brain not as a purely cognitive information processing system, but as a system in which affective functions and cognitive ones are inextricably integrated with one another

the extension of cognitive theory to explain and exploit the role of affect in learning is in its infancy

At the same time there have been developments on the more applied level of thinking about learning, particularly in K–12 education. The first fall-out from the computer presence were strongly on the 'cognitive' side of the split. Even when depersonalised drill and practice in basic skills gave way to the idea of 'intelligent tutors', the theoretical emphasis was on diagnosing and correcting errors in reasoning or factual knowledge. However, the past two decades have seen a powerful growth of educational thinking that emphasise the importance of the relationship of the learner to the process

and the content of learning. This movement has supported and been supported by a growing tendency to view the computer as a ‘motivational’ and even a ‘psychological’ presence rather than a logical analytic engine. There is a need for theories that deal with the interplay of these aspects of learning and a body of experience to support the development of such theories.

‘when we change our emotional states, we’re switching between different ways to think’

This paper is driven by a vision of advancing towards a unified theory that will bridge these gaps. Among the authors of this article there are strong differences about what such a theory might look like. Stronger than these differences, however, is our shared goal of redressing the imbalance between affect and cognition, and between theory and practice, to bring balance to the science of learning and to its technologies. The greatest advancements in the science of learning will require the engagement of multiple perspectives. This paper is a first step, bringing together diverse viewpoints and areas of expertise from a growing community, to begin to construct a science of affective learning.

This overview paper is organised into three broad areas:

- building tools and technologies that elicit, sense, communicate, measure, and respond appropriately to affective factors,
- building new models and learning systems that incorporate affect, as a foundation for both new approaches to education and more effective machine learning,
- developing affectively evocative materials, things-to-learn, and learning environments.

Interwoven closely with these technology construction efforts is a growing research effort to develop and refine theories and terminology related to affect in learning.

2. Challenges in affective learning

Scientific findings over the past decade have started to lay the foundation for a better understanding of the role of affect in learning. Research has demonstrated, for example, that a slight positive mood does not just make you feel a little better but also induces a different kind of thinking, characterised by a tendency toward greater creativity and flexibility in problem solving, as well as more efficiency and thoroughness in decision making. These effects have been found among many groups of different ages and professions [3]. The influences on cognition are not limited to positive mood — affective states such as fear, anger, sadness, and joy show up in the brain as different patterns of blood flow, providing one possible explanation for how affect influences brain activity (e.g. Lane et al [4] and Damasio et al [5]). In the case of positive affect, a theory of two separate but interacting dopamine systems has

been proposed for mediating some of the effects positive affect has on cognition [6]. There is also some indication that positive affect increases intrinsic motivation [7].

Although the work in this area is only beginning to be launched, it already suggests that a positive mood is not best for all kinds of thinking, but that certain affective states facilitate some kinds of thinking better than others. Learning research has long recognised the importance of facilitating different ways of thinking — with beliefs such as ‘you don’t understand something unless you understand it in many ways’. In his forthcoming book, *The Emotion Machine*, Marvin Minsky argues, ‘... when we change what we call our ‘emotional states’, we’re switching between different ways to think’ [8].

Among educators and educational researchers, there is a growing recognition that interest and active participation are important factors in the learning process (e.g. Bransford et al [9]). But acceptance of these ideas is based largely on intuition and generalised references to constructivist theorists [10—12]. There is need for new types of studies on the role of affect in learning. We believe that new technologies can play a particularly important role in these efforts, helping us to measure, model, study, and support the affective dimension of learning in ways that were not previously possible.

3. Terminology and theories

One of the problems with studying affect is defining what it is — illuminating a better definition of affect and related terms like emotion, motivation, caring, and so forth. Modern research in this area began before the turn of the last century, when Charles Darwin [13] and William James [14] devoted seminal works to describing emotion, anchoring its description in measurable bodily changes and expressions. In the last century many cognitive scientists and psychologists have advanced theories and definitions of emotion, motivation, and other affective phenomena. Nearly a hundred definitions of emotion had been categorised as of 1981 [15] when Don Norman wrote his now classic essay naming emotion as one of the twelve major challenges for cognitive science [2].

Today, the burgeoning literature on affect includes diverse communities such as psychology, cognitive science, neuroscience, engineering, computer science, sociology, philosophy, and medicine, and this has contributed to a similarly diverse understanding of numerous basic terms related to affect — such as ‘emotion’, ‘motivation’, ‘attention’, ‘reward’, and more. One of the challenges we face is the bringing together of theorists and practitioners from different fields in order to refine the language used with respect to affect and learning.

nearly a hundred definitions of emotion had been categorised as of 1981

Although there are dozens of books on various affective phenomena, there is a lack of theory that engages the topic of

affect in learning. Some of the classic works on affect emphasise cognitive and information processing aspects in a way that can be encoded into machine-based rules, and studied in a learning interaction. The most widely adopted of these is the OCC model of emotion [16]; however, this model does not include many of the affective phenomena observed in natural learning situations, such as interest, boredom, or surprise. Csikszentmihályi [17] has emphasised the tendency for a pleasurable state of ‘flow’ to accompany problem solving that is neither too easy nor too challenging, and there have been other scattered attempts to address emotions involved in learning (e.g. Lepper and Chabay [18], Mandler [19] and Kort et al [20, 21]). However, there is still very little understanding as to which emotions are most important in learning, and how they influence learning. To date there is no comprehensive, empirically validated, theory of emotion that addresses learning.

With respect to motivation in learning, there has been much more work and much more progress, illuminating the role of intrinsic versus extrinsic influences, the influence of how pleasurable past learning experiences have been, the feeling of contributing to something that matters and the importance of having an audience that cares, among other factors [22—28]. Related concepts such as self-efficacy also play a critical role [29—33], and students’ beliefs about their efficacy, in turn, influence them emotionally [29, 30]. Several researchers have integrated both affective and cognitive components of goal directed behaviour into motivation theories (e.g. Maehr [34], Dweck [35], Ames and Archer [36], Dweck and Leggett [37], and Elliott and Dweck [38]). These and many other efforts have provided vast insight into human affect; however, in very few cases are the theories at a level suitable for implementation in an interactive machine model.

The need for more precise theory is being driven today by growing efforts to build technologies that interact with learners — motivating, engaging, and assisting them in challenging new ways. In many of these efforts, the systems need programmed representations and strategies that will perform in real-time interaction with a human learner. The designers of these systems turn to human-human interaction, and its literature, as an example to guide their design. Thus, the intelligent tutoring system research community examines successful human tutoring as a source of inspiration for what might be implemented in machine tutoring systems, and finds, for example, that ‘expert human tutors... devote at least as much time and attention to the achievement of affective and emotional goals in tutoring, as they do to the achievement of the sorts of cognitive and informational goal that dominate and characterise traditional computer-based tutors’ [18]. But what do these expert teachers ‘see’ and how do they decide upon a course of action? The theories, where they do exist, tend to focus on a high-level set of observations and practice, which does not directly translate into the level of detail needed to implement these phenomena into machines.

Theories of affect in learning need to be tested and evolved. Our approach is to extend classical armchair observations and thought experiments with the development and use of new technologies that help elicit, sense, measure, communicate, understand, reflect upon, and respond to emotions in learning

situations. Conducting controlled experiments dealing with affect has always been a challenge, and new technologies are needed to make this process easier. The sections below outline several directions with respect to creating such technologies. The technology development both derives from and contributes to that of the theory; indeed, we hope, by engaging in both simultaneously, that one will strengthen the other, bringing both closer to clarity and unity.

4. Enabling technologies that sense and respond

One of the reasons understanding about affect has lagged behind that of cognition is that affective state information is hard to measure. You can easily measure someone’s ability to recall a list of learned items, and with somewhat more difficulty, you can test their ability to generalise and apply some learned information. However, it is much harder to measure how they feel while doing these things. How can various tools of learning, and future robots and environments, objectively sense if a learner is pleased, engaged, disengaged, frustrated or ready to quit? And in what ways can these tools enable reflection and discovery about affect? There is a need to develop sensors and interfaces, together with new signal processing, pattern recognition, and reasoning algorithms for assessing and responding to the affect of the learner in real time.

4.1 Sensing without interfering

Affective experience — such as how much pleasure, frustration, or interest you felt — is typically measured by questionnaire (e.g. Matsubara and Nagamashi [39], de Vicente and Pain [40], and Whitelock and Scanlon [41]). Special instruments have been developed in many cases, such as for evaluating the motivational characteristics of an instructor’s classroom delivery (e.g. Keller and Keller [42]). Despite the convenience and widespread acceptance of questionnaires, the use of self-report information is considered unreliable when it comes to emotion — for adults, self-report is coloured by awareness of internal state, reflections on how such a report will be perceived, ability to articulate what one feels, and more. For children, emotion self-report ‘is never highly valid, and any report before age 11 is unwise’ [43]. On top of these problems, questionnaires require interrupting the learning experience, and thus cannot be used unobtrusively and continuously during learning. Alternatives, such as use of external human observers to label affect, are labour and time-intensive, and do not scale to important visions such as that of the Computing Research Association (CRA) to provide ‘A (Computer) Teacher for Every Learner’ [44]. If a machine is to observe a learner continuously, as a skilled mentor or tutor, then it will need skills of affect perception.

affective state information is hard to measure

With skills of affect perception a computer that detects the learner making a mistake while appearing curious and engaged could leave the learner alone since mistakes can be important for facilitating learning and exploration; however, if

the learner is frowning, fidgeting, and looking around while making the same mistake, then the computer might use this affective feedback to encourage a different strategy.

A number of researchers have raised the concern that you cannot begin to measure or respond to affect until after you articulate a clear theory of affect. While theory and clearer terminology are important goals we share, there are examples from natural systems that suggest we can still forge ahead, despite the state of the theory. For example, dogs presumably have no theory of what affect is and yet they appear to sense and respond to their owner's moods, responses that in many cases bring about beneficial consequences. One can make a related argument for infants, who show an ability to respond to how something is said, long before they understand what is said [45]. Thus, we believe that even without a fully-fledged theory of affect, machines can be given some capabilities to recognise and respond to affect. In fact, it is our experience that efforts to build a phenomenon that is poorly understood will aid in helping improve the understanding of that very phenomenon, so that engaging simultaneously in both the practice and the theory helps advance both.

infants respond to *how* something is said long before they understand *what* is said

Emotion recognition is a component of emotional intelligence [46, 47], and skilled humans can assess emotional signals, in themselves and in others, with varying degrees of accuracy. Recent developments in affective computing aim to also give computers skills of emotional intelligence, including the ability to recognise emotion as well as a person might [48]. The basic approach is to observe a person's patterns of behaviour via sensors such as cameras, microphones, or pressure sensors applied to objects the learner is in contact with (mouse, chair, keyboard, steering wheel, toy), and use computers to associate these patterns with probable affective state information. Thus, a camera and computer, equipped with pattern recognition software, might be used to recognise facial muscle movements associated with a smile, and the smile-detection might then be used to help reason about the probability the person is actually happy. (Expressions do not always imply the existence of underlying feelings.) The job of the computer is to assess a constellation of such patterns and relate them to the user's affective state. The latter is what is termed 'emotion recognition' even though it does not really see what you are feeling, but only a pattern of measurable external changes associated with feelings.

Most prior work on emotion expression recognition from speech, image, and video has focused on deliberately expressed emotions, and not on those that occur in natural situations such as classroom learning. The results make it hard to predict rates we can expect for recognising emotions during learning. In general, people can recognise one of about six different emotional states from speech with about 60% accuracy [49]. Computer algorithms match this accuracy

under more restrictive assumptions, such as when the sentence content is already known. However, automated speech recognition that works at about 90% accuracy on neutrally spoken speech tends to drop to 50–60% accuracy on emotional speech [50]. Improved handling of emotion in speech is important for improving recognition of what was said, as well as how it was said. Facial expression recognition is easier for people, and the rates computers obtain are higher — from 65–98% accuracy on tests to date, with the lower numbers on the more 'natural' data that did not control for head movement. Here, the latest research has focused on recognising specific muscle movements known as 'facial actions' [51] that can be used to construct any facial expression [52–59]. Under certain restricted conditions the automated recognisers have been shown to perform comparably to humans trained in recognising facial actions [52]. Combining visual information with other modalities can give improved results [60–63].

Although the progress in facial, vocal, and combined facial/vocal expression recognition is promising, the numerical results given above are on pre-segmented data of a small set of sometimes exaggerated expressions, or on a small subset of hand-marked singly occurring facial actions. The state of the art in affect recognition is similar to that of speech recognition decades ago when the computer could classify the carefully articulated digits, '0, 1, 2, ..., 9,' spoken with pauses in between, but could not accurately detect these digits in continuous conversations. Moreover, we are interested in computer recognition of truly experienced emotions in learning situations, as opposed to emotions that have been expressed by actors or by subjects posed in front of a camera or microphone. Thus we cannot expect the computer to perform perfectly at recognition, and our methods will have to take into account uncertainty factors.

Recently, a number of projects have tackled the sensing and modelling of emotion in learning and educational gaming environments [20, 64–68]. Systems developed in the MIT Media Laboratory include 'expression glasses,' which discriminate upward facial expressions such as those of interest and openness from downward expressions such as those of confusion or dissatisfaction [69]; these are also designed to hide the wearer's expression from view of the teacher, allowing the students to anonymously communicate expressions of confusion to the teacher in real time without fear of what the teacher might think of the student's intellect. Additionally, the Media Lab has attained from 89–96% classification accuracy of three levels of cognitive-emotional stress [70], although the latter data was from drivers in Boston and not from children in learning situations, where stress would also be interesting to study.

Recently, a system was designed at the Media Lab for automated recognition of a child's interest level in natural learning situations (see Fig 1), using a combination of information from chair pressure patterns sensed using Tekscan pressure arrays (recording how postures moved during learning) [71] and from upper facial features sensed using an IBM BlueEyes video camera (<http://www.almaden.ibm.com/cs/blueeyes>) [64]. New algorithms were developed with the aim of seeing if the machine could match a teacher's ratings of affective labels. Training and testing algorithms on separate



Fig 1 Various sensors can capture postural, facial, skin-surface, and gestural changes that carry affective information. From left to right: chair with Tekscan pressure sensors, BlueEyes camera, Galvactivator skin conductivity sensor, and pressure mouse.

sequences of data from the learning experiences, we developed a system that achieved an accuracy of 76% on affect category recognition from chair pressure patterns, and 88% on nine ‘basic’ postures that were identified as making up the affective behaviours. Both sets of results are conservative, being trained on a small set of data, and tested on children not seen before. The accuracy rates increase to 82% and 98%, respectively, on children who have had portions of their data included as part of the training process. All of these results are highly significant, confirming that there is strong evidence of affective information in the postural moves of the child. These results show that elements of affect can be measured with results significantly higher than random. Such methods need further improvement, especially to integrate facial, postural, and other behavioural information for jointly analysing the state of the learner and increasing accuracy of the inference.

the state of the art in affect recognition is similar to that of speech recognition decades ago

Additionally, we see the need to evolve new sensors for learning environments where the learner is not seated in front of a computer as well as for the traditional keyboard/monitor/mouse environment. Our efforts include extending the Galvactivator (<http://www.media.mit.edu/galvactivator>), a skin-conductivity sensing glove, to communicate wirelessly with a nearby hand-held computer. Skin conductivity gives a measure of psychological arousal, which is a strong predictor for both attention and memory [72]. We have found that students enjoyed learning about how this signal changed with their level of engagement in various learning activities [73], and we believe it will be an enabling tool for exploring affect in new learning environments. Additionally, we have been exploring uses of a newly developed ‘pressure mouse’ device, a mouse augmented with eight pressure pads that indicate ‘how’ the mouse is being handled. An increase in physical pressure applied to a pressure-sensitive mouse has recently been shown to be associated with frustration caused by poor

usability in a computer interface [74]. With collaborators at MediaLabEurope, we are also developing new comfortable wireless physiological sensors and real-time signal processing algorithms, which will be useful for monitoring learner stress.

Finally, in neuroscience itself, there is a pressing need to develop ways to measure the levels of interest and motivation of an animal engaged in tasks, and to gauge the levels of stress experienced by the animal. We have been approached by neuroscientists who would like to collaborate in the development of new affect sensing technologies for fundamental research looking at the mechanisms of motivation and attention in the animal brain. This research, while very basic, is needed to inform the understanding of human brain disorders related to attention and motivation.

4.2 Reflecting and interpreting

The extent to which emotional upsets can interfere with mental life is no news to teachers. Students who are anxious, angry, or depressed do not learn; people who are caught in these states do not take in information efficiently or deal with it well [47].

How can systems that measure affective information help people make sense of what has been measured, and respond in useful, appropriate, and respectful ways? One important response is to help people become more aware of their affect — building a kind of ‘affective mirror’ in which the learner is encouraged to reflect upon how their state is influencing their learning experience. Emotional awareness, in oneself and in others, is considered to be a learnable skill of emotional intelligence. Being aware of one’s state, such as frustration, can be instrumental in helping deal with that state productively.

The Galvactivator, which converts level of skin conductivity to the brightness of a glowing LED, is one device that makes it easy to visualise how your psychological arousal changes as you go about activities. We observed classrooms of students wearing these, where the light glowed brightly when they were engaged in discussing ideas or writing in their journals, and went dim (for many of them) when they were lectured to. Skin conductivity exhibits changes with respect to attention and engagement, and reflects interesting patterns when there are disorders of these, such as in autistics and in patients with

various emotion-related deficits (e.g. Bechara et al [75], Hirstine et al [76]). There are also many possible additional means for people to learn about and reflect upon affective signals such as the one shown in Fig 2.

Another project we have initiated is the use of a physically animated computer for facilitating awareness of affect in learning. This project equips a desktop monitor with the ability to move in subtly expressive ways in response to its user. The physical animation of the machine is inspired by natural human-human interaction — when people work together, they move in reciprocal ways, such as shifting posture at conversational boundaries and leaning forward when interested. The physically animated computer will sense and interpret multimodal cues from the user via sensors such as those above. It will then respond to the user's cues with carefully crafted subtle mechanical movements and occasional auditory feedback, using principles derived from natural human-human interaction.

The initial version of this device will be designed to mirror affect from the user in a way that is non-distracting. For example, if the child's face and posture show signs of intense interest in what is on the screen, the computer would hold very still so as to not distract the child. If the child shifts her posture and moves in such a way that suggests she is taking a break, or starting to become bored, the computer will do similarly. In doing so, the system not only acknowledges the presence of the child and shows respect for her level of attentiveness, but also shows subtle expressions of mirroring that, in human-human interaction, are believed to help build rapport and liking [77]. By increasing likeability, we hope to facilitate task outcomes such as how long the child perseveres with the learning task. Of course, the system can also use subtle movements to reflect other expressions such as frustration or disappointment [78]. We are interested in evaluating the impact of such communication on the learner's reflection of her own state, as well as on other performance characteristics of the learning experience.




Fig 2 Prototype of physically animated computer.

We also see value in integrating new affect sensing, recognition, and reflection technologies into efforts to build intelligent tutoring systems and other automated systems where there is potential to adapt the learning experience based on signs of interest, frustration, and any other affect-related cues. We currently have one such collaboration with the AutoTutor project at Memphis (<http://www.autotutor.org>), and there are a growing number of opportunities in this area related to the CRA vision [44] of a 'teacher for every learner.'

By embedding these technologies in learning interactions with automated systems (animated tutors, robotic computers, etc) and also integrating them into other learning environments (see below) we hope to better answer such questions as: What affective states are most important to learning and how do these states change with various kinds of pedagogy? How does knowledge of one's affective state influence outcomes in the learning experience? Additionally, these technologies form the basis for building systems that will interact with learners in more natural ways, bootstrapping the machine's own ability to learn, the topic of the next section.

5. Machines that learn with you

Machine learning has focused largely on algorithms that can label new data, and not on systems that learn naturally from interacting with you. We wish to shift this focus to enable new kinds of systems that learn with people through natural interaction. A key part of this effort will involve the development of new theories and models for integrating affective and cognitive mechanisms used in learning. In so doing, we hope to realise three equally important goals. Firstly, we wish to advance the state of the art in machine learning to develop systems that can learn far more quickly, more broadly, and continuously from natural human instruction and interaction than they could alone. Secondly, we aspire to achieve a deeper understanding of human learning and development by creating integrated models that permit an in-depth investigation into the social, emotional, behavioural, and cognitive factors that play an important role in human learning. Thirdly, we want to use these models and insights to create engaging technologies that help people learn better.

The history of combining affective mechanisms with cognitive ones for improving machine processing goes back at least to the work of Herb Simon [1], who articulated the construction of motivational and emotional controls over cognition, and proposed incorporating these into information processing systems. His work was inspired by Neisser [79], who, in a criticism of the dominant information processing theories, emphasised these fundamental characteristics of human thought:

- human thinking always takes place in, and contributes to, a cumulative process of growth and development,
- human thinking begins in an intimate association with emotions and feelings which is never entirely lost,
- almost all human activity, including thinking, serves not one but a multiplicity of motives at the same time.

Some have claimed that the emotional components may be, in some ways, dominant; Don Norman wrote, ‘There must be a regulatory system that interacts with the cognitive component. And it may well be that it is the cognitive component that is subservient, evolved primarily for the benefit of the regulatory system, working through the emotions, through affect’ [2]. Thus, these pioneering theorists suggested that cognition, and by extension learning, took place within the context of emotional, motivational, perceptual and behavioural structures that shaped those very processes.

By contrast, the dominant trend in machine learning has been to eschew built-in structure or a priori knowledge of the environment and to discover structure that is in the data or the world through extensive search and/or sophisticated statistical learning techniques. Pattern recognition and reinforcement learning are two problem domains in particular that have attracted attention and met with success. In pattern recognition, the system’s goal is to learn a mapping from a set of input features to an output label. The input features might be associated with a gesture, a face, or an acoustic pattern, for example, and the output label might be something like ‘appears to be happy.’ The system typically learns the mapping through a statistical analysis of hundreds or thousands of training examples chosen by a ‘knowledgeable external supervisor’ [80], in which an example contains both the input features and the desired output label. Typically, the system has no a priori knowledge of the structure of the input space and must discover it based on the examples provided by the supervisor. In the domain of reinforcement learning, the goal of the system is to learn an optimal sequence of actions that will move the system from an arbitrary state to a goal state. The main approach of reinforcement learning is to probabilistically explore states, actions and their outcomes to learn how to act in any given situation.

learning in nature is characterised by fast and robust, albeit constrained, learning

Reinforcement learning is an example of unsupervised learning in that the only supervisory signal is the reward received when it achieves the desired goal. However, as with supervised learning techniques, the actual learning algorithm has no a priori knowledge about the structure of the state and action spaces and must discover any structure that exists on its own through its exhaustive exploration of these spaces. As a result, reinforcement learning typically requires hundreds or thousands of examples in order to learn successfully.

Thus, the progress to date in machine learning has come with some caveats. Firstly, the most powerful techniques rely on the availability of an enormous number of training examples. Secondly, they tend not to be appropriate when the environment is changing so quickly that earlier examples are no longer relevant. Thirdly, the underlying representations used in machine learning typically make it difficult for the

systems to generalise from learning one particular thing or strategy to another type of thing. Fourthly, little attention has been paid to the question of how a human naturally guides and scaffolds the learning process, such as calling attention to the part of the task that matters most. Fifthly, and not insignificantly, few would argue that current approaches to machine learning, however successful, have much to tell us about how learning occurs in animals and humans.

By contrast, any survey of animal learning will quickly convince one that learning in nature is characterised by fast and robust, albeit, constrained learning [81, 82]. For example, a dog can be trained to roll over in response to an arbitrary verbal or gestural cue in as little as 20 to 50 repetitions [83]. A nightingale can learn to imitate the song of another bird after as few as five presentations [84]. A typical child learns an average of 8—10 words a day over their first five years [85]. How is it that animals and children can solve these learning problems so effortlessly?

Our hypothesis is that the answer lies not in finding the ‘magic bullet’ of a unitary learning algorithm or module, but rather in discovering the combination of underlying structures and processes that radically simplify what would otherwise be a complex learning problem. In nature, these internal structures are cognitive, behavioural, social, emotional, motivational, shallow and deep, innate and learned, purposed and repurposed. Indeed, an important way that internal structures simplify the learning task is by acting so as to bias the learner to take maximal advantage of external environmental and social-emotional interactions that serve to structure and constrain the learning task. Hence, learning is the result of a complex interplay of structures and processes, both internal and external to the learner, and having both cognitive and affective aspects.

We are beginning to develop computational models and learning systems that capture these key characteristics. We believe that such models can provide new insights into numerous cognitive-affective mechanisms, and shape the design of learning tools and environments, even if they do not compare to the marvellous nature of those that make children tick. The three co-learning scenarios presented below illustrate the kinds of systems we aim to build — each one learns in partnership with people, but with different emphases.

5.1 *A curious robot*

Imagine a robot that exhibits curiosity (Fig 3). Curiosity is a trait of natural learning systems (i.e. people and animals) that exhibit inquisitiveness and a drive to learn, which tends to be followed by quickly learning what they ought to learn, when they ought to learn it, in an ongoing way. Inspired by nature, we envision a curious robot to be a pro-active learner that seeks out experiences and people from which to learn new things.

Humans are natural and motivated teachers for entities that are rewarding to work with. A curious robot will be able to leverage the rich social nature that is uniquely characteristic of human learning and instruction to constrain and bias its own exploration and discovery of new skills and knowledge, thereby

allowing itself to learn more quickly, broadly, and continuously than it could alone. To do so, a curious robot will need a deeper understanding of the learning process, beyond turning the statistical crank of a learning algorithm, to actually reflect upon the learning process — when to learn (or when to get help to learn), what to learn, from whom, how (e.g. recognise success, correct errors, judge progress), and why. In humans and animals, both cognitive and affective factors play an important role in this process. By building this robot, we would further illuminate these factors.

For instance, there is a need to model cognitive-affective mechanisms of attention [86] and saliency measures to allow the robot to determine the significance of stimuli or events, either on its own or when guided by a person (via gesture or directed gaze). When learning something new, it is also important for the robot to assign affective value to incoming stimuli to help bias what it learns — for example: Was the outcome good or bad, am I making progress toward a desired outcome, does my co-learner appear pleased? Internal cognitive-affective mechanisms help assess the affective value of both internal and external states. For example, social referencing (where an infant looks to the expressive reaction of his or her caregiver) plays an important role in helping an infant (as it could for a robot) to affectively evaluate novel situations and to guide his/her subsequent exploration [87]. This will allow our system to learn quickly from natural social interaction, and allow us to explore cognitive-affective models of saliency and affective value with respect to learning in people and animals [88].

The success of such a robot can be measured by its ability to engage the human's natural interest and attention while learning a variety of new skills, tasks, and knowledge from natural human instruction, without requiring any adjustment of the internal learning mechanisms. We believe this will go a

long way to enabling a new class of technological artefacts that readily adapt and learn within the human environment.

5.2 A teachable interactive character

Next, imagine a scenario in which a child teaches a 3-D computer-animated puppy new tricks. Teachable agents are a new area of research [89, 90] and show promise not only in motivating learners, but also in engaging them in opportunities to reflect on attitudes about learning and other meta-learning concerns. Animal training can be viewed as a coupled system in which the trainer and the animal co-operate so as to guide the animal's exploration to discover how to perform new skills. Animal trainers have developed techniques such as 'luring', 'shaping' and 'clicker training' that allow the person to guide the animal's learning from its observed behaviour alone [91, 92]. Because the trainer cannot see inside the animal, it is very important that the virtual puppy's behaviour be an immediate and accurate reflection of what it has learned so far. Moreover, to the extent it can infer, even minimally, the trainer's intent, and use that knowledge to guide what it learns from the trainer, it will be markedly easier to train, and learn in far fewer examples. Blumberg and his colleagues [93, 94] have developed an autonomous animated dog that can be trained using these very techniques and embodying many of the characteristics that make dogs a tractable animal to train. This will provide the child with immediate and compelling feedback as to the success or failure of his or her teaching efforts.

One of the strong differences between animal learning and machine learning to date is the recognition that, in animals, important behaviours are often self-motivating, and are later repurposed for their ultimate use. Thus, a kitten perfects its pounce, not on its prey, but on its littermates, apparently because they are enjoyable to pounce on. Later, the motivational context dictates the object and form of the




Fig 3 A co-learning system involving cognitive-affective mechanisms that regulate both internal and external processes.

pounce — for pleasure, the pounce may aim for an object that is particularly ‘satisfying’ on which to pounce (and perhaps repeatedly so), whereas when hungry, the cat chooses the object and form of pounce that maximises the likelihood of satisfying hunger [95]. Indeed, Panksepp suggests that most so-called appetitive behaviours are self-rewarding and provides a sketch of the possible neural underpinnings of such a mechanism [96]. We would like to construct self-motivating behaviours in machines to illuminate how such mechanisms contribute to learning.

This new technology would also provide children with an opportunity to learn about the importance of motivation as well as context, timing, exploration, and built-in biases for efficient learning. By putting the child in the position of helping to guide the puppy’s learning, we hope to encourage the child to ask questions about his or her own learning, and consider how human learning and teaching may be similar but also different from animal training in interesting ways.

Such a system will have succeeded if the child gains a deeper understanding of the teaching-learning process and uses it to adapt his or her behaviour and hypotheses about learning and teaching.

5.3 *A learning companion*

Finally imagine a scenario in which a machine serves as a computerised learning companion to facilitate a child’s own efforts at learning. A learning companion will not be an intelligent tutoring system that already knows the answers about the subject being learned, but rather a player on the side of the student — a collaborator of sorts — there to help the child learn, and in so doing, learn how to learn better. To do so, the companion will help to keep the child’s exploration going, by occasionally prompting with questions or feedback, and by watching and responding to aspects of the affective state of the child — watching especially for signs of frustration and boredom that may precede quitting, for signs of curiosity or interest that tend to indicate active exploration, and for signs of enjoyment and mastery, which might indicate a successful learning experience. It will have succeeded if students, especially those who encounter frustration and routinely handle it by quitting, learn instead how to persevere, increasing their ability and desire to engage in self-propelled learning.

A computerised learning companion allows for controlled explorations of communicative factors such as the role of facial expression, empathy, mirroring postures, and even emotional contagion, all of which can play a role in human-human interaction and relationship development [97, 98]. While most people cannot bring all these movements under precise control, a computational agent such as the learning companion can [99]. This is not to say that agent-synthesised movements can exactly replace those of people, nor that such theories developed in this environment will exactly map to the human-human environment; however, this level of control does allow for careful testing of hypotheses such as: ‘Can a computer companion’s displays of enthusiasm for a topic infect a student in a way similar to that in which a human companion’s displays of enthusiasm can?’

Going a bit further, it is known that the presence of someone who cares, or at least appears to care, can be motivating [100]. Various studies have linked interpersonal relationships between teachers and students to motivational outcomes over the long term [101—103]. Although computers do not ‘care’ in the sense of having feelings like people have, it is nonetheless possible for them to model certain behaviours and give some of the other impressions that contribute to a perception of caring, as has been recently demonstrated by a ‘relational’ agent built to interact with people over a long period of time [104]. While that research was applied to health (exercise behaviour change), the findings that relate to caring and motivation may be similar to those that influence learning. Although we do not expect that machine ‘caring’ could provide any kind of real substitute for genuine human caring, we do hypothesise that certain aspects of it could be given to learning tools and technologies in ways that have a positive impact upon learners.

6. Fostering love of learning

There may seem to be a dramatic difference in the modernity of two major branches of our research on affective learning. While inanimate objects that react to affect (discussed above) were inconceivable until very recently, people have always created evocative objects that elicit delight, desire, or fear and these objects, whether their makers knew it or not, profoundly affected the way people learn. What is new to the digital age is the scientific understanding and the technology needed to continue this process in a deliberate theory-based manner. We present our strivings in this direction by exploring relationships between technology and learning on four levels each more complex than the previous and further away from the dominant paradigms.

6.1 *Level 1 — the holding power of the computer*

It is widely recognised that computer technology can generate great intensity of engagement. In many cases, the dominant effect comes from dramatic graphics and dynamic colourful animations with no intrinsic connection with the intellectual content. More interesting cases make use of the capacity of the computer to provide a controllable level of challenge in facing problems that are connected with the content. Csikszentmihályi has highlighted the dynamic between ‘challenge’ and ‘mastery’ in the learning process [17, 105]. Too often, educators try to make things ‘easy’ or ‘entertaining’ for learners and seem to think that children need to be ‘motivated’ to do harder work.

a computerised learning companion allows for controlled explorations of communicative factors

A view of activities as motivating in themselves has been theorised by Csikszentmihályi who has found that people become most deeply engaged in activities that are challenging, but not overwhelming. Similarly, Papert has developed the concept of ‘hard fun’ — learners do not mind,

and benefit most from, activities that are ‘hard’ as long as they connect deeply with their interests and passions [106].

6.2 *Level 2 — making it personal*

Less widely studied than intensity of engagement is quality of engagement. A long search for conditions that favour quality of engagement has led us to give special weight to the following four areas.

6.2.1 Constructionism — the engagement of the builder

The simplest description of ‘constructionist’ activities [106, 107] emphasises a cognitive aspect — learners construct new knowledge most effectively when they are in the process of constructing something external which they can examine for themselves and discuss with others. But we also recognise affective aspects. Learners feel differently about the knowledge when they experience themselves as active participants with control over (and personal involvement in) the learning process. And the way they feel about the knowledge profoundly influences what they will do with it and especially how they reflect on it, which in turn influences how it grows and connects.

6.2.2 The physical and the digital

The constructionist principle can, of course, be effective when the constructs are virtual entities. But it has a special quality when the constructs combine the digital and the physical as in the case of the ‘programmable brick’ (developed in the Media Lab and commercialised as LEGO Mindstorms), which allows children to build (among other things) simple robots and endow them with behaviours. Here the physical nature of the construct allows the children to draw on their sophisticated skills and intuitions for sensing and manipulating the environments in which they live while the digital programmability allows them to turn these intuitions into formal knowledge.

But this description in terms of ‘knowledge’ leaves out vital dimensions. The children’s emotional attachment to the objects they have known, their likes and dislikes, their aesthetic judgements all come into play. Pioneers in early childhood education, particularly Froebel and Montessori, attached importance not only to the conceptual but also the relational and aesthetic aspects of objects they designed for children. Indeed, Frank Lloyd Wright credited his boyhood experiences with Froebel’s gifts (the manipulative materials developed for the first kindergarten in early 19th century) as the foundation of his architecture [108].

Of course learning through attachment to objects can benefit learning science as well as art. Even a superficial eye can see that learners are more engaged when they learn principles of physics and engineering by building functioning machines. Our research has shown that this engagement comes, in large part, from personal ‘identification’ with the robots and machines that they build. In our *Beyond Black Boxes* project, we found that students, by building their own robotic constructions made stronger (as well as clearer) connections with the scientific concepts underlying their investigations [109].

This is in line with findings by a growing number of researchers (e.g. Lave and Wenger [110]) who have argued that people form their strongest relationships with knowledge through concrete representations and activities — very different from the formal, abstract representations and approaches favoured in traditional school curricula (particularly in the domains of math and science). The physical-digital combination vastly expands the range of knowledge that can be experienced in this affect-supported fashion by a process that can be called ‘making the abstract concrete’ [111].

6.2.3 Bodies of knowledge

When children program their LEGO Mindstorms constructs they draw on knowledge of many parts of the physical world. The affective force is greatest when this part of the physical world happens to be one’s own body. This observation suggests designing anthropomorphisable constructs, such as the Logo Turtle.

The turtle exists in three forms — as a physical toy that moves on the floor, as a computational object that moves on the screen, and most abstractly as a mathematical entity that plays the role that a point plays in Euclid’s development of geometry [112—114]. With this new representation, children learn important geometric ideas in a more ‘body-syntonic’ way, imagining themselves as the turtle as it draws out geometric patterns, and thus leveraging their intuitions and experiences of their own bodies into more formal knowledge and into a more personal relationship with mathematics.

Numerous other Media Lab projects over the years have contributed to expanding the range of ways in which the body can be ‘morphed’ into mathematics. Knot-tying, piano playing, juggling, skiing, and dance are example domains in which the body-in-motion can support intuitive, emotionally engaging learning about apparently unrelated but potentially deeply interconnected conceptual realms [115—122]. In summer 2003, Media Lab researchers worked with celebrated dancer/educator Jacques d’Amboise on organising a ‘RoBallet workshop’ [123—125] (Fig 4) to explore new directions opened by giving children direct bodily control over affect-rich features of the stage environment — lighting, sound and a projected stage set. RoBallet can be appreciated for a cognitive side and for an affective side: the former about whether planning and thinking about one’s own movements through 3-D space can provide a foundation for thinking like a geometer; the latter about whether doing this in the emotionally intense context of dance gives the learning a special depth and robustness.

6.2.4 Music

Dance brings together two of the potentially richest intersections of the physical and the digital to support affectively powerful learning — movement and music. The latter is being richly pursued in its own right in the Toy Symphony Project [126] (see Fig 5) — the latest step in a tradition of using digital media to give children freer and more direct access to music as a means of creative expression. Everything that has been said about the other dimensions could be repeated here. Instead we mention a new source of insight from the study of musical activities that will surely eventually extend to the others. This is the beginnings of



Fig 4 A RoBallet workshop.

breakthrough in finding the neurological basis of the relationship between affect and learning [127]. It is too soon to read clear conclusions from such studies. But like the brain studies on 'being in love' [128] they offer a glimpse into a possibility that the study of cross-modality learning such as music or dance with mathematics or language might be the route to understanding how the differences in affect are actually based in differences on neurological function.




Fig 5 Hyperscore extract.

This area, together with those discussed in the three preceding sections, involves bringing the physical body into the learning experience.

Exploring how the role of the body interacts with cognition is a growing area of interest in psychology and cognitive science (e.g. Bausalon et al [129]), where new theories are being developed to explain many of the powerful intertwined influences of affective, cognitive and other bodily systems.

6.3 Level 3 — affective epistemology

6.3.1 Knowing how, knowing that, and getting to know you

On the previous levels we explored relationships between knowledge and action and between different bodies of knowledge. On this level we talk about giving knowledge new forms.

We begin by noting that the learning we have just imagined does not fit into the clean separation between 'knowledge' and 'affect' that lies behind such statements as 'l'affectivité constitue le ressort des actions ... et en règle l'énergie.' (Affect is the spring of actions ... and governs their energy) [10]. Piaget's statement calls up an image of a child motivated by the prospect of a reward (affect) to memorise the fact (propositional knowledge) that a circle is such and such or to learn how (procedural knowledge) to draw one. Something different is happening here. The child is getting to know and perhaps to like the mediating technology as one might get to know and like a person. Papert makes much in his book *Mindstorms* of the fact that the kind of learning involved in getting to know a person is not reducible to propositional or procedural knowledge [113]. Nobody would question that getting to know a person engages affect in deep and essential ways. It is not about facts and skills. It is about relationship.

The most important learning is not merely energised by affect; it is affective, and forms a relationship with the learner. What the child internalises is less like what logicians study than like what psychoanalysts talk about as introjecting a person. The facts of geometry are not learned as such; the child has acquired the ability to re-generate them by internalising an entity with which he can identify enough to pilot through geometric manoeuvres by connecting to and drawing upon his/her own bodily knowledge.

Before turning to an example of what this kind of theorising can mean in practice we note that it points to a set of theoretical sources for educational thinking whose common characteristic is leading to a view of 'the stuff' of which mind is composed as being more like an active creature than like the propositions and links in passive data structures. On the computational side it points to theories in the style of *The Society of Mind* [130] and the kind of thinking that lies behind object-oriented programming [131]. On the psychological side it points to psychoanalytic theory and most especially by its 'object-relations' branch that gives a central role to the internalisation of 'objects' (including, in fact especially, people [132, 133]. In each case the 'stuff of mind' is not something like a proposition or a procedure but more like an active being.

6.3.2 Making mathematics that people will love to learn

The example mentioned above is a thread of research inspired by the slogan: 'Instead of trying to make children love the math they hate, make a math they'll love' [112]. This programme¹ raises two kinds of issue related to affect that are essentially different from those discussed earlier in this paper.

¹ Mathematics here is a placeholder for all areas of study. We focus on one for the sake of concreteness.

Firstly, a difference of time-scale — most of the earlier examples concern transitory emotional states such as being happy (or interested or bored) which can change in seconds, while here we are looking at what are sometimes called dispositions such as ‘being in love’ or ‘being interested in baseball’ which typically have far longer duration. Secondly, a difference in purpose — in the earlier examples, our focus was on understanding or guiding a process of learning, whereas here we put a focus on what is learned, i.e. how can things-to-learn be designed so as to elicit affect in ways that will facilitate learning.

In its scope our project is most like the ‘New Math’ of mid-century which did try to use a learning theory to ‘make a new math.’ But instead of being discouraged by its failure, we offer an explanation that expresses in another way the central position of this paper. The learning theory used there was entirely cognitive — based on considerations of what was logically age-appropriate and ignoring all affective issues. It was about understanding not about loving. It failed not because the idea of changing the content of school mathematics is wrong; it failed because its changes went in exactly the wrong direction. Our successes in limited innovations suggest that the major weakness of traditional school math is being too dissociated from personal feelings and physical applications. The emphasis on logic in the design of the ‘New Math’ aggravated these weaknesses.

The above observations point to the direction in which we are working. Others can be found in a special number of the *International Journal of Computers for Mathematical Learning* [134–136] and earlier publications by principals at the Media Lab [106, 113, 137]. An attempt at putting the whole together can be found in Papert’s paper ‘An Exploration in the Space of Mathematics Educations’ [138].

6.4 *Level 4 — the social side of affective learning*

6.4.1 **Roots, fruits, and shoots**

Learning is rooted in the person and the culture; it bears fruit through the construction process; it has shoots that branch into new areas, shaping and transforming the community around the learner. These principles of learning are experientially based, differing markedly from the concept that requires a disconnected accumulation of chunks of knowledge. In order for the learning to become truly rooted, a person has to have a deep emotional attachment to the subject area. Rooting and the possibilities for branching flow from a better understanding of emotion, motivation, attention, comfort, community, and culture.

Rogoff [139] describes the progression of community learning through the planes of apprenticeship, guided participation, and participatory appropriation. We have been building technological affordances that serve in the role of emotional and inspirational mentors and that foster the creative and idiosyncratic connections to learning that help community members to progress through these planes.

Digital technologies offer new opportunities for discovering roots, adapting to preferences, and enabling creative and idiosyncratic connections to learning and knowing. There is a

need for a new range of expressive technologies and a more integrated methodology to facilitate rooted knowledge construction and support development of shoots to new areas through electronic collaboration and support [140].

An example is found in our work discovering engine culture in Thailand [141]. Numerous local innovations and widespread knowledge made it clear there was a deeply rooted culture of learning and practice building upon knowledge from the internal combustion engine. This became evident in our Project Lighthouse when rural adolescents, all of whom had left school after only a few years, used a variety of computational technologies to design a new dam and address critical water problems in the region. Not only was it remarkable that they learned enough to design an irrigation system without the usual years of formal preparation, but also through their local knowledge and ‘engine culture’ spirit they succeeded where the government had repeatedly failed.

A second example is found in our longstanding efforts to immerse adults and children in the hard, but fruitful, work of inquiry and storytelling [142, 143], where we have seen communities forge around print and radio journalism, creating stories of interest and concern for themselves and their audiences. This act of expression, facilitated by easy-to-use tools, led to an active debate over the content of their stories and, more importantly, the processes that they engage in as media producers. A collaborative editing process seems to help them develop a critical stance towards traditional media. As community participants challenge each other, they begin to understand the biases and critical thought processes that are the norm for professional journalists.

learning is rooted in the person and the culture

A third example flows from our work on new content enabled by computational media. We have engaged learners in developing computational models to improve life in their community [123]. We find that the content finds roots in a wider range of people and thus diminishes equity issues in accomplishment in mathematics and science across gender, class, and racial lines [125]. We have also found that these activities help root knowledge in these domains, make connections to tacit knowledge in the learner, and facilitate branching into new areas within these domains.

6.4.2 **Wear learning**

‘We need to brand math, and all learning, so that each morning, when youth stand in front of the mirror deciding who they will be that day, they always decide to wear learning’ — Christine Ortiz, youth leader of Florida’s ‘Truth’ campaign.

In 1998, a group of Florida teens was given authority and resources to launch a comprehensive campaign to change teen smoking behaviour. The teens established a network of grassroots, youth-led community organisations, and organised a mass-media outreach that included television commercials (both professional and, more effective, ‘unpolished’ or homemade commercials), magazine

advertisements, billboards, and flyers. After the first year of what was branded the 'Truth' campaign, the programme had a 92% recognition rate, equal to mega-brands like Nike. In the first three years of the programme, smoking declined with unprecedented rates across the state.

Affect about learning has a social component. The declining attitude pattern shown in Fig 6 is only partially due to how mathematics is presented in schools; it also reflects a dislike of mathematics that is deeply ingrained in contemporary cultures. The study of affect and learning should include looking at how such cultural affect comes about and how it changes.




Fig 6 Student attitudes about mathematics decline from Grade 8 to Grade 12 (1996). (Graph is from US Department of Education [144])

Can a youth-led campaign, equipped with the right scientific findings about affective learning and the right technologies, help transform feelings about learning and about mathematics as the Truth campaign transformed feelings about smoking? Although it might be unrealistic to expect everyone to develop a deep love of mathematics, there is no reason why everyone could not love learning. We would like to see the trend in these attitude graphs reversed. We accept as an ultimate challenge for 'Affective Learning' the realising of ideas such as this.

Acknowledgements

We are grateful to many colleagues and students in the MIT community who helped shape the ideas expressed here. In particular, we would like to thank Ann Graybiel, Hiroshi Ishii, Barry Kort, Gary McDarby, Marvin Minsky, Nicholas Negroponte, Sile O'Modhrain, and Ted Selker for inspiring discussions and ideas related to this research agenda. We would also like to thank our families for their support of our time and effort invested in this research.

References

- Simon H A: 'Motivational and emotional controls of cognition', in *Models of Thought*, pp 29—38, Yale University Press, New Haven (1967).
- Norman D A: 'Twelve issues for cognitive science', in Norman D A (Ed): 'Perspectives on cognitive science', pp 265—295, Erlbaum, Hillsdale, NJ (1981).
- Isen A M: 'Positive affect and decision making', in Lewis M and Haviland J (Eds): 'Handbook of emotions', Guilford, New York (2000).
- Lane R, Reiman E M, Ahern G L, Schwartz G E and Davidson R J: 'Neuroanatomical correlates of happiness, sadness and disgust', *American Journal of Psychiatry*, 154, pp 926—933 (1997).
- Damasio A R, Grabowski T J, Bechara A, Damasio H, Ponto L L B, Parvizi J and Hichwa R D: 'Subcortical and cortical brain activity during the feeling of self-generated emotions', *Nature Neuroscience*, 3, pp 1049—1056 (2000).
- Ashby F G, Isen A M and Turken U: 'A neuropsychological theory of positive affect and its influence on cognition', *Psychological Review*, 106, No 3, pp 529—550 (1999).
- Estrada C, Isen A M and Young M J: 'Positive affect influences creative problem solving reported source of practice satisfaction in physicians', *Motivation and Emotion*, 18, pp 285—299 (1994).
- Minsky M: 'The Emotional Machine', (2003) — <http://web.media.mit.edu/~minsky/>
- Bransford J D, Brown A L and Cocking R R: 'How people learn: brain, mind, experience and school', National Academy Press, Washington, DC (1999).
- Piaget J: 'Six Etudes de Psychologie', Pub Gonthier, Paris (1971).
- Vygotsky L: 'Thought and Language', The MIT Press, Cambridge, MA (1962).
- Vygotsky L: 'Mind in Society: the development of the higher psychological processes', The Harvard University Press, Cambridge, MA. Originally published 1930, Oxford University Press, New York (1978).
- Darwin C: 'The Expression of Emotions in Man and Animals', The University of Chicago Press, Chicago, IL (reprinted in 1965) (1872).
- James W: 'Writings 1879-1899', Chapter on Emotion, pp 350—365, The Library of America (1992) (originally published in 1890 as *Principles of Psychology*, Holt, NY (1890)).
- Kleinginna Jr, P R and Kleinginna A M: 'A categorised list of emotion definitions, with suggestions for a consensual definition', *Motivation and Emotion*, 5, No 4, pp 345—379 (1981).
- Ortony A, Clore G L and Collins A: 'The Cognitive Structure of Emotions', Cambridge University Press, Cambridge (1988).
- Csikszentmihályi M: 'Flow: The Psychology of Optimal Experience', HarperCollins (1991).
- Lepper M R and Chabay R W: 'Socialising the intelligent tutor: bringing empathy to computer tutors', in Mandl H and Lesgold A (Eds): 'Learning issues for intelligent tutoring systems', pp 242—257 (1988).
- Mandler G: 'Mind and Body: Psychology of Emotion and Stress', W W Norton & Co, New York, NY (1984).
- Kort B, Reilly R and Picard R W: 'An affective model of interplay between emotions and learning: reengineering educational pedagogy-building a learning companion', in *Proceedings of International Conference on Advanced Learning Technologies (ICALT 2001)*, Madison Wisconsin (August 2001).
- Kort B, Reilly R and Picard R W: 'External representation of learning process and domain knowledge: affective state as a determinate of its structure and function', *Workshop on Artificial Intelligence in Education (AI-ED 2001)*, San Antonio, (May 2001).
- Vroom V H: 'Work and Motivation', Wiley, New York (1964).
- Keller J M: 'Motivational design of instruction', in Reigeluth C M (Ed): 'Instructional Design Theories and Models: an overview of their current status', Erlbaum, Hillsdale, NJ (1983).
- Keller E F: 'A feeling for the organism: the life and work of Barbara McClintock', San Francisco, W H Freeman (1983).
- Keller J M: 'Strategies for stimulating the motivation to learn', *Performance and Instruction*, 26, No 8, pp 362—632 (1987).

- 26 Keller J M: 'IMMS: instructional materials motivation survey', Florida State University (1987).
- 27 Ames C: 'Classrooms: goals, structures, and student motivation', *Journal of Educational Psychology*, 84, No 3, pp 261—271 (1992).
- 28 Vail P: 'Emotions: the on/off switch for learning', Modern Learning Press (1994).
- 29 Bandura A: 'Self-efficacy: toward a unifying theory of behaviour change', *Psychological Review*, 84, pp 191—215 (1977).
- 30 Bandura A: 'Social learning theory', Prentice-Hall, Englewood Cliffs, NJ (1977).
- 31 Pajares F: 'Self-efficacy beliefs in academic settings', *Review of Educational Research*, 66, pp 543—578 (1996).
- 32 Schunk D H: 'Self-efficacy and achievement behaviours', *Educational Psychology Review*, 1, pp 173—208 (1989).
- 33 Zimmerman B J: 'Self-efficacy: an essential motive to learn', *Contemporary Educational Psychology*, 25, pp 82—91 (2000).
- 34 Maehr M L: 'Meaning and motivation: toward a theory of personal investment', in Ames R and Ames C (Eds): 'Research on motivation in education', pp 39—73, Academic Press, San Diego, CA (1984).
- 35 Dweck C S: 'Motivational processes affecting learning', *American Psychologist*, 41, pp 1040—1048 (1986).
- 36 Ames C and Archer J: 'Achievement goals in the classroom: students' learning strategies and motivation processes', *Journal of Educational Psychology*, 80, pp 260—267 (1988).
- 37 Dweck C S and Leggett E L: 'A social-cognitive approach to motivation and personality', *Psychological Review*, 95, pp 256—273 (1988).
- 38 Elliott E S and Dweck C S: 'Goals: an approach to motivation and achievement', *Journal of Personality and Social Psychology*, 54, pp 5—12 (1988).
- 39 Matsubara Y and Nagamachi M: 'Motivation systems and motivation models for intelligent tutoring', in Frasson C (Ed): *Proceedings of the Third International Conference in Intelligent Tutoring Systems*, pp 139—147 (1996).
- 40 De Vicente A and Pain H: 'Motivation diagnosis', in *Intelligent Tutoring Systems* (1998).
- 41 Whitelock D and Scanlon E: 'Motivation, media and motion: reviewing a computer supported collaborative learning experience', in Brna P et al (Eds): 'Proceedings of the European Conference', (1996).
- 42 Keller J M and Keller B H: 'Motivational delivery checklist', Florida State University (1989).
- 43 Kagan J: Personal communication (2002).
- 44 Computing Research Association — <http://www.cra.org/>
- 45 Fernald A: 'Intonation and communicative intent in mother's speech to infants: Is the melody the message?' *Child Development*, 60, pp 1497—1510 (1989).
- 46 Salovey P and Mayer J D: 'Emotional intelligence', *Imagination, Cognition and Personality*, 9, No 3, pp 185—211 (1990).
- 47 Goleman D: 'Emotional Intelligence', Bantam Books, NY (1995).
- 48 Picard R W: 'Affective Computing', The MIT Press, Cambridge, MA (1997).
- 49 Scherer K R: 'Chapter 10: speech and emotional states', in Darby J K (Ed): 'Speech evaluation in psychiatry', Grune and Stratton Inc (1981).
- 50 Hansen J: 'Speech under stress', ICASSP'99, Phoenix, Arizona (1999).
- 51 Ekman P and Friesen W: 'Facial action coding system', Consulting Psychologists Press (1977).
- 52 Cohn J F, Zlochower A J, Lien J and Kanade T: 'Automated face analysis by feature point tracking has high concurrent validity with manual FACS coding', *Psychophysiology*, 36, pp 35—43 (1999).
- 53 Bartlett M S, Hager J C, Ekman P and Sejnowski T J: 'Measuring facial expressions by computer image analysis', *Psychophysiology*, 36, pp 253—263 (1999).
- 54 Donato G, Bartlett S, Hager M, Ekman J C P and Seinowski T: 'Classifying facial actions', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21, No 10, pp 974—989 (1999).
- 55 Tian Y, Kanade T and Cohn J F: 'Recognising action units for facial expression analysis', *IEEE Trans Pattern Analysis and Machine Intelligence*, 23, No 2 (February 2001).
- 56 Kapoor A and Picard R W: 'Real-time, fully automatic upper facial feature tracking', in *Proceedings of the 5th International Conference on Automatic Face and Gesture Recognition 2002*, Washington DC (May 2002).
- 57 Kapoor A and Picard R W: 'A real-time head nod and shake detector', *Workshop on Perceptive User Interfaces*, Orlando, FL (November 2001).
- 58 Mota S and Picard R W: 'Automated posture analysis for detecting learner's interest level', *1st IEEE Workshop on Computer Vision and Pattern Recognition, CVPR HCI 2003* (2003).
- 59 Bartlett M S, Littlewort G, Braathen B, Sejnowski T J and Movellan J R: 'A prototype for automatic recognition of spontaneous facial actions', in Becker S and Obermayer K (Eds): 'Advances in Neural Information Processing Systems', The MIT Press, Cambridge, MA (2003).
- 60 DeSilva L C, Miyasato T and Nakatsu R: 'Facial emotion recognition using multi-model information', *Proceedings of IEEE International Conference on Information, Communication and Signal Processing*, Singapore, pp 397—401 (1997).
- 61 Huang T S, Chen L S and Tao H: 'Bimodal emotion recognition by man and machine', *ATR Workshop on Virtual Communication Environments*, Kyoto, Japan (April 1998).
- 62 Chen L S, Huang T S, Miyasato T and Nakatsu R: 'Multimodal human emotion/expression recognition', in *Proceedings of International Conference on Automatic Face and Gesture Recognition*, Nara, Japan, IEEE Computer Society (April 1998).
- 63 Kapoor A, Picard R W and Ivanov Y: 'Probabilistic combination of multiple modalities to detect interest', *Proceedings of Int Conf on Pattern Recognition*, Cambridge, England (August 2004).
- 64 Kapoor A, Mota S and Picard R W: 'Towards a learning companion that recognises affect', *AAAI Fall Symposium 2001*, North Falmouth, MA (November 2001).
- 65 Conati C: 'Probabilistic assessment of users' emotions in educational games', *Applied Artificial Intelligence*, 16, Nos 7 and 8, pp 555—575 (2002).
- 66 Zhou X and Conati C: 'Inferring user goals from personality and behaviour in a causal model of user affect', *Intelligent User Interfaces 2003*, pp 211—218 (2003).
- 67 Zhou X and Conati C: 'Modeling students' emotions from cognitive appraisal in educational games', *Intelligent Tutoring Systems 2002*, pp 944—954 (2002).
- 68 Sheldon-Biddle E, Malone L and McBride D: 'Objective measurement of student affect to optimise automated instruction', *Proceedings of Workshop on Modelling User Attitudes and Affect, User Modeling '03* (2003).
- 69 Scheirer J, Fernandez R and Picard R W: 'Expression glasses: a wearable device for facial expression recognition', *CHI'99 Short Papers*, Pittsburgh, PA (1999).
- 70 Healey J and Picard R W: 'SmartCar: detecting driver stress', *Proceedings of ICPR'00*, Barcelona, Spain (May 2000).
- 71 Tekscan: 'Tekscan Body pressure measurement system user's manual', Tekscan Inc, South Boston, MA USA (1997).
- 72 Reeves B and Nass C: 'The Media Equation', Center for the Study of Language and Information, Cambridge University Press (1996).

- 73 Picard R W and Scheirer J: 'The Galvactivator: a glove that senses and communicates skin conductivity', Proceedings from the 9th International Conference on Human-Computer Interaction, New Orleans (August 2001).
- 74 Dennerlein J T, Becker T, Johnson T, Reynolds C and Picard R W: 'Frustrating computer users increases exposure to physical risk factors', Proceedings of International Ergonomics Association, Seoul, Korea (2003).
- 75 Bechara A, Damasio H, Tranel D and Damasio A: 'Deciding advantageously before knowing the advantageous strategy', *Science*, 275, pp 1293—1295 (1997).
- 76 Hirstine W, Iversen P and Ramachandran V S: 'Autonomic responses of autistic children to people and objects', Proceedings of the Royal Society, London, B, 268, pp 1883—1888 (April 2001).
- 77 La France M: 'Posture Mirroring and Rapport', in David M (Ed): 'Interaction rhythms: periodicity in communicative behaviour', Human Sciences Press, New York, NY, pp 279—298 (1982).
- 78 Liu K and Picard R W: 'Subtle expressivity in a robotic computer', presented at CHI 2003 Workshop on Subtle Expressiveness in Characters and Robots (2003).
- 79 Neisser U: 'The imitation of man by machine', *Science*, 139, pp 193—197 (1963).
- 80 Sutton R and Barto A: 'Reinforcement learning: an introduction', The MIT Press, Cambridge, MA (1998).
- 81 Shuttleworth S J: 'Cognition, evolution and behaviour', Oxford University Press, New York, NY (1998).
- 82 Gould J and Gould C: 'The animal mind', W H Freeman, New York, NY (1999).
- 83 Wilkes G: 'Click and treat training kit', Click and Treat Inc, Mesa, AZ (1995).
- 84 Marler P: 'Song learning: the interface between behaviour and neuroethology: philosophical transactions', *Biological Sciences*, 329, No 1253, pp 109—114 (1990).
- 85 Markham E: 'Constraints children place on word meanings', *Cognitive Science*, 14, pp 57—77 (1990).
- 86 Breazeal C, Edsinger A, Fitzpatrick P and Scassellati B: 'Active vision systems for sociable robots', in Dautenhan K (Ed): 'IEEE Transactions on Systems, Man and Cybernetics', Part A, 31, No 5, pp 443—453 (2001).
- 87 Feinman S, Roberts D, Hsieh K, Sawyer D and Swanson D: 'A critical review of social referencing in infancy', in Feinman S (Ed): 'Social referencing and the social construction of reality in infancy', pp 15—54, Plenum Press, New York (1992).
- 88 Breazeal C, Hoffman G and Lockerd A: 'Teaching and working with robots as a collaboration', in Proceedings of AAMAS 2004, Columbia NY (2004).
- 89 Biswas G, Katzlberger T, Bransford J, Schwartz D and TAG-V (the teachable agents group at Vanderbilt): 'Extending Intelligent Learning Environments with Teachable Agents to Enhance Learning', 10th International Conference on AI in Education: AI-ED in the wired and wireless future (Moore J D, Redfield C L and Johnson W L (Eds)), IOS Press, Amsterdam, pp 389—397 (2001).
- 90 Biswas G, Schwartz D, Bransford J and TAG-V (the teachable agents group at Vanderbilt): 'Technology Support for Complex Problem Solving: From SAD Environments to AI, Smart Machines in Education', in Forbus K D and Feltovich P J (Eds): 'The coming revolution in education technology', AAAI/MIT Press, Menlo Park, CA (2001).
- 91 Pryor K: 'Clicker Training for Dogs', Sunshine Books Inc, Waltham, MA (1999).
- 92 Lindsay S R: 'Applied Dog Behaviour and Training', Iowa State University Press, Ames, IA (2000).
- 93 Blumberg B and Downie M: 'Integrated learning for interactive synthetic characters', *Transactions on Graphic*, 21, No 3, Proceedings of ACM SIGGRAPH 2002 (2002).
- 94 Blumberg B: 'D-learning: what learning in dogs tells us about building characters that learn what they ought to learn', in Lakemeyer G and Nebel B (Eds): 'Exploring Artificial Intelligence in the New Millenium', Morgan Kaufman Publishers, San Francisco, CA (2002).
- 95 Lorenz K and Leyhausen P: 'Motivation of human and animal behaviour: an ethological view', Van Nostrand Reinhold Co, New York, NY (1973).
- 96 Panksepp J: 'Affective Neuroscience', Oxford University Press (1998).
- 97 Bull P E: 'Posture and Gesture', International Series in Experimental Social Psychology, Pergamon Press (1987).
- 98 Hatfield E, Cacioppo J and Rapson R L: 'Emotional contagion', New York, Cambridge University Press (1994).
- 99 Burleson W, Picard R W, Perlin K and Lippincott J: 'A platform for affective agent research', in Proceedings of the AAMAS Workshop on Empathetic Agents, Columbia University, New York, NY (July 2004).
- 100 Wentzel K: 'Student motivation in middle school: the role of perceived pedagogical caring', *Journal of Educational Psychology*, 89, No 3, pp 411—419 (1997).
- 101 Pianta R C: 'Beyond the parent: the role of other adults in children's lives', *New Directions in Child Development*, 57, Jossey-Bass, San Francisco (1992).
- 102 Wentzel K and Asher S R: 'Academic lives of neglected, rejected, popular, and controversial children', *Child Development*, 66, pp 754—763 (1995).
- 103 Birch S H and Ladd G W: 'Interpersonal relationships in the school environment and children's early school adjustment: the role of teachers and peers', in Juvonen J and Wentzel K (Eds): 'Social motivation: understanding children's school adjustment', Cambridge University Press, New York (1996).
- 104 Bickmore T and Picard R W: 'Establishing and maintaining long-term human-computer relationships', *Transactions on Computer Human Interaction*, to appear (2004).
- 105 Csikszentmihályi M: 'Finding Flow: the psychology of engagement with everyday life', HarperCollins (1997).
- 106 Papert S: 'The Children's Machine', Basic Books (1993).
- 107 Kafai Y and Resnick M (Eds): 'Constructionism in practice: designing, thinking, and learning in a digital world', Mahwah, NJ, Lawrence Erlbaum (1996).
- 108 Brosterman N: 'Inventing Kindergarten', Harry Abrams Publishing (1997).
- 109 Resnick M, Berg R and Eisenberg M: 'Beyond black boxes: bringing transparency and aesthetics back to scientific investigation', *Journal of the Learning Sciences*, 9, No 1, pp 7—30 (2000).
- 110 Lave J and Wenger E: 'Situated Learning: legitimate peripheral participation', University Press, Cambridge (1991).
- 111 Turkle S and Papert S: 'Epistemological pluralism', *Signs*, 16, No 1, pp 128—157 (1990).
- 112 Papert S: 'Teaching children to be mathematicians versus teaching about mathematics', *International Journal of Mathematics Education in Science and Technology*, 3, pp 249—262 (1972).
- 113 Papert S: 'Mindstorms: Children, Computers, and Powerful Ideas', Basic Books (1980).

- 114 Abelson H and diSessa A: 'Turtle Geometry', The MIT Press (1980).
- 115 Austin H: 'A computation theory of physical skill', MIT Artificial Intelligence Laboratory Note AIM-330 (1974).
- 116 Burton R et al: 'Skiing as a Model of Instruction', in Rogoff B and Lave J (Eds): 'Everyday Cognition: development in social context', Harvard University Press (1984).
- 117 O'Modhrain S: 'Playing by feel: incorporating haptic feedback into computer-based musical instruments', unpublished PhD Dissertation, Stanford University (2000).
- 118 Sloane N J A and Wyner A D (Eds): 'Claude Elwood Shannon: Collected Papers', Piscataway, NJ, IEEE Press (1940/1993).
- 119 Strohecker C: 'Why knot?', unpublished doctoral dissertation, MIT Media Laboratory, Cambridge, MA (1991).
- 120 Strohecker C: 'Learning about topology and learning about learning', Proceedings of the Second International Conference on the Learning Sciences, Association for the Advancement of Computing in Education (1996).
- 121 Strohecker C: 'Understanding topological relationships through comparisons of similar knots', *AI&Society: Learning with Artifacts*, 10, pp 58—69 (1996).
- 122 Strohecker C: 'Cognitive zoom: from object to path and back again', *Spatial Cognition II*, pp 1—15 (2000).
- 123 Cavallo D, Blikstein P, Sipitakiat A, Basu A, Camargo A, de Deus Lopes R and Cavallo A: 'The city that we want: generative themes, constructionist technologies and school/social change', in Proceedings of IEEE Technology and Education in Developing Countries 2004, Joensuu, Finland (2004).
- 124 Cavallo D, Sipitakiat A, Basu A, Bryant S, Welti-Santos L, Maloney J, Chen S, Asmussen E, Solomon C and Ackermann E: 'RoBallet: exploring learning through expression in the arts through constructing in a technologically immersive environment', in Proceedings of International Conference on the Learning Sciences 2004, Los Angeles, California (2004).
- 125 Cavallo D, Sipitakiat A, Basu A and Bryant S: 'Opening pathways to higher education through engineering projects', in Proceedings of the American Society of Engineering Education 2004, Salt Lake City, Utah (2004).
- 126 Machover T: 'Toy symphony', Boosey & Hawkes, New York and London (2002/3) — <http://www.toysymphony.net>
- 127 Tomaino C: 'How music can reach the silenced brain', *Cerebrum: The Dana Forum on Brain Science*, 4, No 1, pp 23—24 (2002).
- 128 Lim M M et al: 'Ventral striatopallidal oxytocin and vasopressin V1a receptors in the monogamous prairie vole', *J Comp Neurol*, 468, No 4, pp 555—570 (January 2004).
- 129 Barsalou L W, Niedenthal P M, Barbey A K and Ruppert J A: 'Social embodiment, in the psychology of learning and motivation: advances in research theory', Ross B H (Ed), 43, Academic Press, Elsevier Science, USA (2003).
- 130 Minsky M: 'The Society of Mind', Simon & Schuster, New York, NY (1985).
- 131 Kay A: 'Computer software', *Scientific American*, 251, No 3, pp 41—47 (September 1984).
- 132 Winnicott D W: 'Playing and reality', London, Tavistock (1971).
- 133 Greenberg J R and Mitchell S A: 'Object Relations in Psycho-Analytic Theory', Harvard Press (1983).
- 134 Resnick M: 'Thinking Like a Tree (and Other Forms of Ecological Thinking)', *International Journal of Computers for Mathematical Learning*, 8, No 1, pp 43—62 (2003).
- 135 Wilenski U: 'Statistical mechanics for secondary school: The GasLab Modeling Toolkit', *International Journal of Computers for Mathematical Learning*, 8, No 1, pp 1—41 (2003).
- 136 Papert S: 'Foreword', Special Issue on Multi-Agent Programming, *International Journal of Computers for Mathematical Learning*, 8, No 1, (2003).
- 137 Resnick M: 'Turtles, Termites and Traffic Jams', The MIT Press (1994).
- 138 Rogoff B: 'Observing sociocultural activity on three planes: participatory appropriation, guided participation and apprenticeship', in Wertsch J V, Rio P D and Alvarez A (Eds): 'Sociocultural Studies of the Mind', Cambridge University Press, Cambridge, pp 139—164 (1995).
- 139 Papert S: 'An exploration in the space of mathematics educations', *International Journal of Computers for Mathematical Learning*, 1, No 1, pp 95—123 (1996).
- 140 Cavallo D: 'Epistemology all the way down', EuroLogo 2003, Porto, Portugal (2003).
- 141 Cavallo D: 'Emergent design and leaning environments: building on indigenous knowledge', *IBM Systems Journal*, 39, Nos 3 and 4, pp 768—781 (2000).
- 142 Smith B, Bender W, Driscoll J, Endter I, Turpeinen M and Quan D: 'Silver stringers and junior journalists: active information producers', *IBM Systems Journal*, 39, Nos 3 and 4 2000.
- 143 Monroy-Gomez C: 'eRadio,' MIT SM Thesis (2004).
- 144 US Department of Education, National Center for Education Statistics, The Condition of Education 2000, NCES 2000-62, US Government Printing Office, Washington, DC (2000).



Rosalind (Roz) Picard is co-director of the Media Lab's Things That Think consortium and head of the Lab's Affective Computing group, and author of the award-winning book, *Affective Computing* (MIT Press). She is also the author of over 90 peer-reviewed scientific articles on varied topics, and has served as an editor and guest editor for several prestigious publications; she has also been featured in national and international forums for the general public, such as *The New York Times*, *Nightline*, the *BBC*, and *Vogue*. She earned a BS in electrical engineering from the Georgia Institute of Technology and was named a National Science Foundation Graduate Fellow. She earned her MS and PhD, both in electrical engineering and computer science, from MIT.



Seymour Papert is a mathematician and one of the early pioneers of artificial intelligence. In addition, he is internationally recognised as the seminal thinker regarding computers and pedagogy for children. His collaboration with Jean Piaget at the University of Geneva led him to consider using mathematics in the service of understanding how children can learn and think. In the early 1960s, he came to MIT, where, with Marvin Minsky, he founded the Artificial Intelligence Lab and co-authored their seminal work, *Perceptrons* (1970). He is also the author of *Mindstorms: Children, Computers and Powerful Ideas* (1980), and *The Children's Machine: Rethinking School in the Age of the Computer* (1992). He has written numerous articles about mathematics, artificial intelligence, education, learning, and thinking.



Walter Bender is executive director of the MIT Media Laboratory, a senior research scientist, director of the Electronic Publishing group, and a member of the Laboratory's Information Organised consortium. He also directs the Gray Matters special interest group, which focuses on technology's impact on the aging population. A founding member of the Media Lab, he studies new information technologies, building upon the interactive styles associated with existing media and extending them into domains where a computer is incorporated into the

interaction. He has participated in much of the pioneering research in the field of electronic publishing and personalised interactive multimedia. He received his BA from Harvard University, and his MS from MIT.



Tod Machover is a composer and the head of the Media Lab's Hyperinstruments/Opera of the Future group. His compositions have been praised for breaking traditional artistic and cultural boundaries. In 1995 he was named a "Chevalier de l'ordre des Arts et des Lettres", one of France's highest cultural honours. He has composed five operas, including the celebrated Brain Opera, and is the inventor of Hyperinstruments, a technology using smart computers to augment virtuosity. He is also the creator of the Toy Symphony, an international music performance and education project. He was formerly director of musical research at Pierre Boulez's IRCAM institute in Paris. He received both his BA and MA from the Juilliard School in New York.



Bruce Blumberg is a research scientist at the Media Lab, focusing on building computer systems that can robustly learn the kinds of things animals learn easily, that behave with the everyday commonsense that animals display, and that evoke the feelings of companionship that animals such as dogs invoke in us. He received his PhD in media arts and sciences from MIT, studying with Pattie Maes. He came to the Media Lab from NeXT Inc, where he was one of the original employees, and prior to NeXT was the product manager for the original Apple

LaserWriter. He has an MS from MIT's Sloan School of Management and a BA from Amherst College.



Mitchel Resnick explores how new technologies both necessitate and facilitate deep changes in the ways people think and learn. His Lifelong Kindergarten group has developed a variety of educational tools, including the 'programmable bricks' that were the basis for LEGO's award-winning MindStorms robotics construction kit.

He also led the development of StarLogo, a software toolkit for modelling decentralised systems. He is co-founder and principal investigator for the Lab's Digital Nations consortium. He also co-founded the Computer Clubhouse project, a network of after-school learning centres for youth from underserved communities. He earned a BS in physics from Princeton University, and an MS and PhD in computer science from MIT. Before pursuing his graduate degrees, he worked as a science/technology journalist for Business Week magazine. In 1994, he was awarded a National Science Foundation Young Investigator Award.



Cynthia Breazeal directs the Lab's Robotic Life group and holds the LG Career Development chair, having previously been a postdoctoral associate at MIT's Artificial Intelligence (AI) Lab. She is interested in developing creature-like technologies that exhibit social commonsense and engage people in familiar human terms. Kismet, her anthropomorphic robotic head, has been featured in international media and is the subject of her book *Designing Sociable Robots* (MIT Press). Breazeal earned ScD and MS degrees at MIT in electrical

engineering and computer science, and a BS in electrical and computer engineering from the University of California, Santa Barbara.



Deb Roy heads the Media Lab's Cognitive Machines group, which develops machines that learn to communicate in human-like ways. This work combines elements of spoken language understanding and generation, machine learning, computational linguistics, computer vision, robotics, and developmental psychology. Other active research interests include multilingual technologies, multimedia search and retrieval algorithms, and interactive robotics. He holds a Bachelor of Applied Science in computer engineering from the University of Waterloo, Canada, and an MS and PhD from MIT's Program in Media Arts and Sciences.



David Cavallo co-directs the Lab's Future of Learning group, focusing on the design and implementation of reforms in learning environments and educational systems, on the role that technology can play in this process, and on the design of new technologies for learning. Prior to the Media Lab, he led the design and implementation of medical informatics at Harvard University Health Services, as well as designing and building knowledge-based systems for industry, most notably a set of intelligent microworlds for training air traffic controllers. In addition, he

founded and led the Advanced Technology group for Digital's Latin American and Caribbean region. Cavallo received his MS and PhD from the Program in Media Arts and Sciences at MIT, and a BS in computer science from Rutgers University.



Carol Strohecker directs the Everyday Learning research group at Media Lab Europe. She is concerned with how people think and learn, and how objects, artefacts, technologies and environments can elicit and support these processes. Prior to joining MLE, she worked at MERL (Mitsubishi Electric Research Laboratories) and in the Human Interface Group of Sun Microsystems. She earned a PhD in media arts and sciences from MIT, and an MS in visual studies, also from MIT. She was a Fellow of the Harvard University Graduate School of Design, the US National Endowment for the Arts, and the Massachusetts Council for the Arts and Humanities. She holds four US patents for her work in interactive media tools and methods.