

# Affective Computing

R. W. Picard

MIT Media Laboratory; Perceptual Computing; 20 Ames St., Cambridge, MA 02139  
picard@media.mit.edu, <http://www.media.mit.edu/~picard/>

## Abstract

Recent neurological studies indicate that the role of emotion in human cognition is essential; emotions are not a luxury. Instead, emotions play a critical role in rational decision-making, in perception, in human interaction, and in human intelligence. These facts, combined with abilities computers are acquiring in expressing and recognizing affect, open new areas for research. This paper defines key issues in “affective computing,” computing that relates to, arises from, or deliberately influences emotions. New models are suggested for computer recognition of human emotion, and both theoretical and practical applications are described for learning, human-computer interaction, perceptual information retrieval, creative arts and entertainment, human health, and machine intelligence. Significant potential advances in emotion and cognition theory hinge on the development of affective computing, especially in the form of wearable computers. This paper establishes challenges and future directions for this emerging field.

## 1 Fear, Emotion, and Science

Nothing in life is to be feared. It is only to be understood. – Marie Curie

Emotions have a stigma in science; they are believed to be inherently non-scientific. Scientific principles are derived from rational thought, logical arguments, testable hypotheses, and repeatable experiments. There is room alongside science for “non-interfering” emotions such as those involved in curiosity, frustration, and the pleasure of discovery. In fact, much scientific research funded by defense budgets has been essentially prompted by fear. Nonetheless, emotions are generally regarded as wreaking havoc on reasoning. Although emotions pervade science, their role has been marginalized.

Why bring emotion or affect into any of the deliberate tools of science? Moreover, shouldn’t emotion be completely avoided when considering properties to associate with computers? After all, computers control significant parts of our lives – the phone system, the stock market, nuclear power plants, airplane flights, and more. Who wants a computer to be able to “feel angry” at them? To feel contempt for any living thing?

In this paper I will set forth key issues in what I call “affective computing,” computing that relates to, arises from, or deliberately influences emotions. I will elaborate further on this definition and its implications below.

The topic of emotion is a difficult one to treat scientifically, but that is precisely what needs to be done. In this paper I will illustrate ways in which affective computing can break new

ground in the scientific study of emotions. I will suggest computational models for affect recognition, and will describe new applications of affective computing to areas such as computer-assisted learning, perceptual information retrieval, creative arts and entertainment, and human health. Affective computing is a critical new research area in need of exploration, one which can significantly contribute to advances in emotion and cognition theory, while greatly improving the interaction between humans and machines.

I should state some things that I do not intend “affective computing” to address. I am *not* proposing the pursuit of computerized cingulotomies<sup>1</sup> or even the business of building “emotional computers” in the negative sense of the word “emotional” which implies a loss of desirable rationality. However, I will discuss creative and unpredictable computers.

This paper will also not review the massive literature on emotion and cognition theory; I have only included references where needed to support claims related to affective computing.

I will also not propose answers to the difficult and intriguing questions, “what are emotions?” “what causes them?” and “why do we have them?”<sup>2</sup> It is my hope and expectation that research in affective computing, by using computers to recognize and synthesize emotions, can assist scientists in getting closer to the answers of these important questions.

This paper is organized as follows. In the remainder of this section I briefly describe two recent discoveries to support the importance of emotion in cognition, first in perception, and second in rational decision making. I also present a scenario, based on the role of emotions in learning, as an example of affective computing. Section 2 outlines the key issues to developing computers that can recognize human emotion and express emotion. Section 3 poses human affect recognition as a pattern recognition problem, and proposes models for its solution; this section may be skipped by those who are not engaged in the details of building these systems. Computers that “have” emotion present new moral and ethical dilemmas which are broached in Sect. 4. Computers which can recognize and express affect lead to a myriad of new applications; I have suggested over fifty in this paper, with most appearing in Sect. 5. This research agenda and the contributions of this paper are briefly summarized in Sect. 6.

### 1.1 Songs vs. laws

Let me write the songs of a nation; I don’t care who writes its laws. – Andrew Fletcher

Emotion pulls the levers of our lives, whether it is love that

---

<sup>1</sup>The making of small wounds in the ridge of the limbic system known as the cingulate gyrus, a surgical procedure to aid severely depressed patients.

<sup>2</sup>For a list of twelve open questions in the theory of emotion, see Lazarus [1].

leads to forgiveness, or curiosity that drives scientific inquiry. As humans, our behavior is greatly influenced by the “songs” in our hearts. Rehabilitation counselors, pastors, parents, and politicians know that it is not laws that exert the greatest influence on people but rather the drumbeat to which they march. For example, the death penalty has not lowered the murder rate in the states where it has been instituted as law. However, murder rates are significantly influenced by culture, or the cultural “tune.” Consider the following scenario:

*Imagine that your colleague keeps you waiting for a highly important engagement to which you thought you were both committed. You wait with reason, and with increasing puzzlement by his unusual tardiness. You think of promises his delay is causing you to break, except for the promise you made to wait for him. Perhaps you swear off future promises like these. He is completely unreachable; you think what you will say to him about his irresponsibility. But you still wait, because you gave him your word. You wait with growing impatience and frustration. Maybe you waver between wondering “is he ok?” and feeling so irritated that you say under your breath, “I’ll kill him when he gets here.”*

When he finally shows, after you have nearly given up your last promise of waiting, how do you respond? Whether you greet him with rage or relief, doesn’t his expression affect your response? Your response changes dramatically if he arrives inconsiderately carefree, or with woeful countenance. A mere expression of affect can powerfully influence subsequent behavior.

In saying that emotions, or “songs,” pull the levers of our lives, I am *not* suggesting that laws are unimportant, or even that we do away with the law-based artificial intelligence (AI) systems given to computers, even though the latter are widely acknowledged to be brittle and unintelligent. Rather, I am saying that laws and rules are not necessarily the most important part in human behavior and intelligence. In fact, laws and rules do not even play a solo in such cognitive tasks as perception or decision-making, according to recent neurological evidence. Let’s consider briefly these two activities, beginning with some evidence regarding perception, as illustrated in the next scenario.

## 1.2 Limbic perception

“Oh, dear,” he said, slurping a spoonful, “there are not enough points on the chicken.” – Michael Watson [2].

Synesthetes may feel shapes on their palms as they taste, or see colors as they hear music. Synesthetic experiences behave as if the senses are cross-wired, as if there are no walls between what is seen, felt, touched, smelled, and tasted. However, the neurological explanation for this heightened perceptual phenomenon is not “crossed-wires.”

The neurologist Cytowic has studied the neurophysiology of synesthetic experience [3]. Because the cortex is typically regarded as the home of sensory perception, it is expected to show increased activity during synesthetic experiences, where a person experiences external and involuntary sensations somewhat like a cross-wiring of the senses – for example certain sounds may elicit seeing strong colors. One would expect that during this heightened sensory experience, there would be an increase in cortical activity, perhaps in the parietal lobe’s tertiary association area where the three senses of vision, touch, and hearing converge. However, Cytowic found that scans of cortical blood flow <sup>3</sup> during synesthesia episodes indicate a collapse of cortical

metabolism. An overall increase of brain metabolism occurred, but it was not in the “higher” cortex, where it was expected.

Cytowic’s studies point to a corresponding increase in activity in the limbic system, a collection of parts of the brain which lie predominately between the brain stem and the two hemispheres of the cortex. The limbic system<sup>4</sup> has traditionally been assumed to play a less influential role in perception than the cortex, which lies “above” it. The limbic system is the seat of emotion, memory, and attention. Its activity during synesthesia indicates that the limbic system plays a significant role in perception.

In a recent treatise on emotion theory, Izard [4] describes emotion as both a motivating and guiding force in perception and attention. One does not need a blood-flow scan or theory of emotion, however, to recognize that emotion greatly influences perception. We are all familiar with emotion’s influence on perception from observing this influence in other people – influences that have received names such as the fear-induced “tunnel vision,” or the joy-induced “seeing through rose-colored glasses.”

### 1.2.1 The limbic-cortical tangle

Note that my distinction between cortical and limbic functions is for emphasis only; in practice, normal limbic and cortical brain areas do not operate in isolation, but are functionally intertwined. Not only have the two areas been artificially separated in most studies, but there is a tendency among scientists to attribute higher functionality to the cortex, which is physically higher and much easier to probe.

However, discoveries such as that of the limbic role in the higher function of perception imply that a high or dominating function is not necessarily cortical. Along with the synesthesia findings mentioned above, the research of LeDoux, another neuroscientist, suggests that the hippocampus, long considered the key structure of the limbic system, is significantly involved in registering and making sense of perceptual patterns. LeDoux has also recently mapped the paths in the brain which demonstrate how the limbic system responds to certain events before the cortex is involved, notifying it after other “emergency” measures have been taken. Cytowic points out in his books that there are substantially more connections from the limbic system to the cortex than vice-versa. Some scientists have suggested that these discoveries indicate that the limbic influence may be the greater. Note that this suggestion does not imply we are “run by our passions” as might be spoken of someone who does not act reasonably; rather, it implies that even reasonable behavior is neurologically directed by these so-called passions.

As is often the case with scientific claims, philosophers and artists have anticipated them ahead of scientific researchers. These neurology findings are no exception; philosophers have argued the predominance of emotion for centuries, and artists have incorporated it into their foundational beliefs about aesthetic goals.

Although the role of emotions is powerful; we often overlook their influence. However, consider that we often hear a person say “Sorry, I wasn’t thinking” but not “Sorry, I wasn’t feeling.” Whatever our perception of its role, the limbic system is a crucial player in our mental activity. It is hard to say conclusively which system of the brain is *directing* the show – but if the limbic system is not directing, then it is, at least, an actor that

---

<sup>4</sup>There is not complete agreement on what parts of the brain constitute the limbic system, but it includes at least the hypothalamus, hippocampus in the temporal lobe, and the amygdala.

<sup>3</sup>Measured by the Oberist-Ketty xenon technique.

has won the heart of the audience.

### 1.3 Re-evaluating decision making

It is not surprising that emotions play a critical role in perception; but, what about emotion in more “rational” cognitive functions, such as decision-making? We all know emotions can *influence* decisions, often negatively, but do they play a more significant role? Below I will describe recent neurological evidence that indicates a more surprising, and significant, role for emotions.

#### 1.3.1 The thinking–feeling axis

“Scientific conclusions must be decided with the head; whom you choose to marry may be decided with the heart.” – folk advice

Most people consider that both head and heart are useful for decision-making, as long as they are kept in their rightful place, as in the folk advice above. In fact, people often polarize thoughts and feelings, speaking of them as separate phenomena.

The popular Myers-Briggs personality-type indicator, has “thinking” and “feeling” as opposite endpoints of one of its axes for describing personality. In fact, the Myers-Briggs type indicator reveals a gender bias along this axis, indicating that two-thirds of men tend to lie closer to the “thinking” side and two-thirds of women tend to lie closer to the “feeling” side [5]. This bias sometimes appears in male-female stereotypes, and many books have appeared on its implications for human interaction. Although I do not wish to pursue the male-female distinctions further here, it is worth noting that such differences might also extend to human-computer interaction. As such, affective computers might tend to be considered more “feminine.” However, this is not an appropriate conclusion, given the neurological evidence that both male and female brains rely on emotion in normal thinking. As such, affective computers should not be considered more “feminine,” but more “human.”

Neurologically, no polarization, or clean dividing line occurs between thinking and emotions. In fact, we will find something completely unexpected. Recall that the brain does not cleanly separate cortical and limbic activity:

Authorities in neuroanatomy have confirmed that the hippocampus is a point where everything converges. All sensory inputs, external and visceral, must pass through the emotional limbic brain before being redistributed to the cortex for analysis, after which they return to the limbic system for a determination of whether the highly-transformed, multi-sensory input is salient or not. [6].

The limbic brain is the “home base” of emotion, but it is not the only part of the brain engaged in the experience of emotion. Extensive research by Damasio and his colleagues has identified several non-limbic regions which affect emotion. These findings have been recently summarized in the provocative book [7]. But there is a much bigger surprise in his findings.

#### 1.3.2 Too little emotion impairs decision-making

We all know that too much emotion can wreak havoc on reasoning, but now there is evidence that *too little* emotion can also wreak havoc. This evidence requires a shift from the usual paradigm of how people separate emotions and rationality. I refer the reader to the careful arguments and references collected by Damasio [7] for the justification such a far-reaching paradigm-shift demands, and here provide but a brief explanation of the findings to support the need for affective computers.

Damasio’s patients have frontal-lobe disorders, affecting a part of the cortex that communicates with the limbic system. Otherwise, the patients appear to be intelligent, and *unusually rational*. However, these same patients suffer from an impaired ability to make decisions. Years of studies with frontal-lobe patients indicate that they spend inordinate amounts of time trying to make decisions that those without frontal-lobe damage can make quite easily [7]. For example, the mere task of choosing a date to schedule an appointment can lead these patients through abnormally long chains of decisions, perhaps without ever reaching a decision, until a date is imposed upon them by someone who is tired of waiting for their response.

The frontal-lobe disorder in these patients interferes with their ability to combine emotional limbic responses with their otherwise cortical decision-making. Damasio’s hypothesis is that emotion plays a biasing role in decision-making – establishing the values used in evaluating potential outcomes, and essentially warding off an infinite logical search.

Damasio’s findings support independent scientific arguments for the essential role of emotion. Johnson-Laird and Shafir have recently reminded the cognition community of the inability of logic to determine which of an infinite number of possible conclusions are sensible to draw, given a set of premises [8]. Consider: how do you decide which path to take given some evidence? There is not time to consider *every* possible logical constraint and associated path. Emotion does not merely play a tie-breaking role in making certain decisions; rather, it appears to be essential in learning the biases required to construct rational responses.

Damasio’s findings provide neurological support that there is no “pure reason” in the healthy human brain – emotions are vital for healthy rational human thinking and behavior [7]. His patients are abnormally rational, not too unlike the rule-based programs that comprise today’s models of decision-making.

It must be emphasized at this point that by no means should anyone conclude that logic or reason are irrelevant; they are as essential as the “laws” described earlier. Additionally, the neurological evidence describes an essential role for emotions, the “songs.”

Therefore, these findings indicate that scientific study of emotion should not be merely an interesting side-area of study in cognitive science, but rather, the study of emotion is essential if we are to understand human cognition. The implications are significant also for computer science and industry: computers, if they are to be truly effective at decision-making, will have to have emotion-like mechanisms working in concert with their rule-based systems. “Pure reason” may continue as a Platonic ideal, but in successful cognitive systems, it is a logical howler.

### 1.4 Tests of thinking, tests of feeling

In normal human cognition, thinking and feeling are mutually present. If one wishes to design a device that “thinks” in the sense of mimicking a human brain, then must it both think and feel?

Consider briefly the classic test of a thinking machine: the Turing test.<sup>5</sup> The Turing test examines if, in a typical conversation between two participants who have no sensory contact with each other, the tester cannot tell if the replies are being generated by a human or a machine. Although the test cannot *prove* that a machine does or does not think, it is a terrific exercise in thinking about thinking.

---

<sup>5</sup>With slight modifications from the original proposed by Turing [9].

The Turing test is considered a test of whether or not a machine can think in the truest sense of duplicating mental activity. Since mental activity involves a close coupling between cortical and limbic activity, a test of true thinking must involve a test of emotion.

Consider that one might converse with the computer passionately about a song or a poem, or describe to it the most tragic of accidents. To pass the test, the computer responses should be indistinguishable from human responses. Because humans almost always respond with emotion to highly emotional events, either to empathize or to counter balance, the intelligent computer would need to be capable of recognizing emotion and providing affective responses.

Although the Turing test is usually performed with text-only communication, so that sensory expression, viz., voice intonation and facial expression, do not play a role, emotions are still communicated through the written word. This power and importance of influencing emotion through language was a primary tenet of Aristotle's *Rhetoric* [10]. A machine, even limited to text communication, will be a more effective communicator if given the ability to perceive and express emotions.

Of course the crux of the Turing test is what comprises the questions. Hofstadter has suggested that "humor, especially emotion," would comprise the acid test of intelligence for a "thinking machine" [11]. But the argument for emotion as necessary for intelligence goes far beyond its interplay with humor. Goleman has recently argued that emotion is a hallmark of human intelligence [12], and that "emotional intelligence" can be more important for predicting success in life than traditional IQ tests. Emotional intelligence involves factors such as empathy, which, in a machine, would require at least a keen ability to recognize, understand, and express emotions, if not also the ability to "have" emotions.

## 1.5 Affective communication

An increasing number of people spend more time directly communicating with their computer than with other people. Daily interaction between humans and computers has tremendous psychological impact, not to mention billions of dollars of economic impact. It is not my purpose to review the research in this area, which is covered in numerous conferences; however, I would like to describe one intriguing set of recent results which support the importance of affective computers.

This particular set of studies was recently conducted by Nass and his colleagues at Stanford [13]. They performed a number of classical studies of human social interaction, substituting computers into a role usually occupied by humans. Hence, a test that would ordinarily study a human-human interaction is used to study a human-computer interaction. For example, one experiment might study how what is said by human A about human B's performance changes when A gives the evaluation face-to-face with B, vs. when A gives the evaluation about B to another (presumably neutral) person. In general, humans are nicer face-to-face. In Nass *et al.*'s variation, human B is replaced with computer B, and human A gives an evaluation both to computer B and to another computer. Despite the switch, Nass and colleagues found that the human results still held, e.g. the tendency to be nicer "face to face" still held. Numerous other experiments were done in this vein, revealing that the classic results of the human-human studies were maintained in the human-computer studies. After accounting for potential biasing factors, Nass *et al.* concluded that individuals' interactions with computers are inherently natural and social [13], [14].

Because affective communication occurs naturally between people, it is expected by people when they interact with computers. In fact, we often see people attribute emotion to things that clearly do not have emotion – a wind-up dog that wags its tail, for example. Although people *know* that wind-up toys and computers do not have emotions, nonetheless, their discourse and actions often assume them.

Emotion plays an essential role in communication – even in its subtlest form, where it merely indicates that communication has succeeded, that we are understood. If you reprimand someone and their expression does not change, then the natural inclination is to continue your communication until you receive a visible or verbal sign that it has succeeded. This type of human expectation may perhaps be at the root of the practice of many computer users, who keep typing the same wrong thing at the computer, hoping it will eventually respond differently.

Affect recognition and expression is also necessary for sympathy and communication of understanding, the latter of which is considered one of man's greatest psychological needs [15]. Negroponte, in *Being Digital*, reminds us that even a puppy can tell when you are angry with it [16].

Basic affect recognition and expression is expected by humans in communication. Computer-based communication to date has largely removed or ignored affective bits. A quantum leap in communication will occur when computers become able to recognize and express affect.

## 1.6 Example: The effective and affective piano teacher

Before moving to the key issues and research challenges in affective computing, let's consider an example of its use. One of the interests in the Media Lab is the building of better piano-teaching computer systems; in particular, systems that can grade some aspects of a student's expressive timing, dynamics, phrasing, etc. [17]. This goal contains many challenges, one of the hardest of which involves expression recognition, distilling the essential pitches of the music from its expression. Recognizing and interpreting affect in musical expression is important, and I'll return to it again later. But first, let's consider a scenario, where you are receiving piano lessons from a personal computer teacher:

Imagine you are seated with your computer piano teacher, and suppose that it not only reads your gestural input, your timing and phrasing, but that it can also read your emotional state. In other words, it not only interprets your musical expression, but also your facial expression and perhaps other physical changes corresponding to your emotional feelings. Assume it has the ability to distinguish the three emotions we all appear to have at birth – distress, interest, and pleasure [18],<sup>6</sup>

Given affect recognition, the computer teacher might find you are doing well with the music and you are pleased with your progress. "Am I holding your interest?" it would consider. In the affirmative, it might nudge you with more challenging exercises. If it detects your frustration and many errors, it may slow things down and give you encouraging suggestions. Detecting user distress, without the user making mechanical playing errors, might signify a moving requiem, a sticky piano key, or the need to prompt for more information.

The piano-teacher scenario raises the issue of observing not just someone's emotional expression, but also their underlying emotional state. How do we detect a person's emotions? Is

---

<sup>6</sup>This view of [18] is not unchallenged; facial expression in the womb and on newborns has no broadly accepted explanation.

it via some metaphysical sixth sense? Whether or not such a sense might exist or play a role exceeds the scientific scope of this paper; consequently, this possibility will not be further addressed. However, there are *measurable* ways that appear to significantly influence how we discern emotion – we can measure digitally the signals of voice, facial expression, and other aspects of so-called body language. Moreover, there is evidence that we can build systems that begin to identify both emotional expression, and its generating state. First, however, it is necessary to understand more about how emotions are generated and expressed.

## 2 Physical and Cognitive Aspects of Emotion

What signals assist you in guessing another person’s emotional state? Are there universal physiological measurements from which a computer could discern mood? Are emotions purely “cognitive”, purely “physical,” or some kind of mix? If an emotion causes an accompanying physiological state, can that state be induced to cause that emotion?

The literature on emotion and cognition is still debating most of these questions; in fact, emotion theorists still do not agree even on a definition of emotion. The goal of this section is to provide some background from the literature which I find important to the development of affective computing.

Particularly relevant are the two (generally opposing) treatments of emotion in the literature – emotion as cognition, and emotion as physiological response. Below I will emphasize the role of *both* physical and cognitive components for affective computing. I will also clarify some terminology, and highlight ways in which emotion is both expressed and induced.

It is beyond the scope of this paper to overview the literature with its many theories of emotion; I will refer the reader instead to the collections gathered by Plutchik and Kellerman [19] and to the references at the end of this paper which themselves contain many excellent surveys. The focus below is on the background most relevant to giving computers the ability to recognize, express, and model affect.

### 2.1 Poker face, poker body?

The level of control involved in perfecting one’s “poker face” to hide emotions is praised by western culture. But, can we perfect a “poker body?” Despite her insistence of confidence, you hear fear in her voice; although he refuses to cry in your office, you see his eyes twitching to hold back a potential flood. You spot the lilt in my walk today and therefore expect I am in a good mood. Although you might successfully conceal nervousness in your voice, you may not be able to suppress it throughout your body; its evidence may sneak into a clammy handshake or emerge as cold feet.

Although we cannot observe directly what someone feels (or thinks, for that matter), and they may try to persuade us to believe they are feeling a certain way, we are not easily deceived. Beethoven, after he became deaf, wrote in his conversation books that he could judge from the performer’s facial expression whether or not the performer was interpreting his music in the right spirit [20].

Despite the fact that we are not all experts at reading faces, and comedians and actors can excel at feigning emotions, it is claimed that the attentive observer is always able to recognize a false smile [21].<sup>7</sup> This is consistent with the findings of

<sup>7</sup>This view is debated, e.g., by [1] who claims that all phenomena that change with emotion also change for other reasons,

Duchenne over a century ago:

The muscle that produces this depression on the lower eyelid does not obey the will; it is only brought into play by a genuine feeling, by an agreeable emotion. Its inertia in smiling unmask a false friend. [21]

Neurological studies also indicate that emotions travel their own special path to the motor system. If the neurologist asks a patient who is paralyzed on one side to smile, then only one side of the patient’s mouth raises. But when the neurologist cracks a funny joke, then a natural two-sided smile appears [22]. For facial expression, it is widely accepted in the neurological literature that the will and the emotions control separate paths:

If the lesion is in the pyramidal system, the patients cannot smile deliberately but will do so when they feel happy. Lesions in the nonpyramidal areas produce the reverse pattern; patients can smile on request, but will not smile when they feel a positive emotion. – Paul Ekman in [21].

In other words, a faked smile travels a different path than a genuine one. Not only does this imply that, physiologically, false and sincere smiles may be discriminated, but it illustrates the existence of multiple paths for emotional expression. When I talk later about how computers might be given emotions, it is important to remember that emotion-generating mechanisms distinct from ordinary rule-based systems might be required.

To give computers the ability to recognize emotions, we will need to consider how the internal emotional state is communicated physiologically, i.e., in ways that are externally observable. Debate persists about the nature of the coupling between emotion and physiological response, although most scientists now accept both a physiological and a cognitive component in their treatment of emotion. In the following sections I will discuss key issues related to the physiological components, then the cognitive components,<sup>8</sup> and finally, inducement between the physical and cognitive.

### 2.2 Terminology

Before proceeding, it is helpful to clarify terminology. *Sentic*<sup>9</sup>, *emotional*, and *affective* are used interchangeably in this paper, although I will tend to use “sentic” to refer to the physical manifestations of emotion. An affective *state* refers to your internal dynamic state when you have an emotion. This emotional state cannot be directly observed by another person, but may be inferred.

All you consciously perceive of your own emotional state is referred to as your emotional *experience*. Some authors equate emotional experience with emotional “feelings” but I prefer to minimize use of the word “feelings” as it can be ambiguous with sensory feelings, e.g. feeling a pinprick.

What you reveal to others, voluntarily (perhaps even a false smile) or not (what you do not succeed in suppressing), is your emotional *expression*. Expression via the motor system, or “sentic modulation” is usually involuntary, and is one clue which others observe to guess your emotional state.

but these claims are unproven.

<sup>8</sup>Of course the cognitive components, i.e., the brain, are also physical. Treating the body and mind separately can lead to errors, as captured by the title of [7]. The separation will be used here primarily to distinguish arguments from the literature.

<sup>9</sup>“Sentic” is from the Latin *sentire*, the root of the words “sentiment” and “sensation” [23].

Finally, “mood” tends to refer to a longer-term emotional state, although duration might be difficult to quantify given that moods can “swing” abruptly.

### 2.3 Physiological aspects of emotion: sentic modulation

There is a class of qualities which is inherently linked to the motor system ... it is because of this inherent link to the motor system that this class of qualities can be communicated. This class of qualities is referred to commonly as emotions.

In each mode, the emotional character is expressed by a specific subtle modulation of the motor action involved which corresponds precisely to the demands of the sentic state.

– Manfred Clynes [23]

The body usually responds physically to an emotion, although James’s 1890 view of this response *being* the emotion is not accepted today. Nonetheless, the motor system acts as a carrier for communicating emotional state, what I call “sentic modulation” after the foundational principles Clynes established in this area [23].

Sentic modulation (e.g. voice inflection, facial expression, posture) is how an emotional state is typically expressed; it is the primary means of communicating human emotion. When computers learn to recognize human emotion, they will rely primarily on sentic modulation. To give computers affect recognition requires understanding the physical manifestations of emotion.

A number of emotion and cognition theorists have studied the physiological correlates of emotions. Lazarus *et al.* [24] argue that each emotion probably has its own unique somatic response pattern, and cite other theorists who argue that each has its own set of unique facial muscle movement patterns.

Facial expressions are one of the two most widely acknowledged forms of sentic modulation. Duchenne, in his 1862 thesis (republished in [21]) identified independent expressive face muscles, such as the muscle of attention, muscle of lust, muscle of disdain or doubt, and muscle of joy. Most present attempts to automate recognition of facial expression are based on the subsequent Facial Action Coding System of psychologist Paul Ekman [25], which provides mappings between muscles and an emotion space.

The second widely acknowledged form of sentic modulation is via voice intonation: you can hear love in her voice, anxiety in his. Vocal emotions can be understood by young children before they can understand what is being said [26] and by dogs, which we assume can not understand what is being said. Voice, of course, is why the phone has so much more bandwidth than email or a written letter. Spoken communication transcends the message of the words.

Other forms of sentic modulation have been explored by Clynes in his pioneering book, *Sentics* [23]. One of his principles, that of “sentic equivalence,” allows one to select an arbitrary motor output of sufficient degrees of freedom for the measurement of “essentic form,” a precise spatiotemporal dynamic form produced and sensed by the nervous system, which carries the emotional message. The form has a clear beginning and end, that can be expressed by various motor outputs. By this principle, emotional state could be determined from outputs other than facial expression or voice.

The motor output explored most carefully by Clynes is the transient pressure of a finger during emotional expression. In

these experiments the subject deliberately expresses an emotional state by pressing against a measuring surface while experiencing that state. The finger-pressure response has been measured for thousands of people, and found to be not only repeatable for an individual, but to reveal distinct traces for states such as no emotion, anger, hate, grief, love, joy, sex, and reverence [23] across groups of individuals, and to some extent, cultures. Clynes suggests that these traces are indicative of the underlying essentic form. Other forms of motor output such as chin pressure (for a patient who was paralyzed from the neck down) and foot pressure have yielded comparable characteristic essentic forms.

There are many physiological responses which vary with time and which might potentially be combined to assist in recognition of emotional states. These include heart rate, diastolic and systolic blood pressure, pulse, pupillary dilation, respiration, skin conductance and color, and temperature. These forms of sentic modulation will be revisited near the end of this paper in the discussion on affective wearable computers.

Given the human is experiencing an emotion, e.g. hate, then certain values of motor system observations such as a tense voice, glaring expression, or finger pressure strongly away from the body are most probable. Respiration rate and heart rate may also increase. In contrast, given feelings of joy, the voice might go up in pitch, the face reveal a smile, and the finger pressure have a slight bounce-like character. Even the more difficult-to-analyze “self-conscious” emotions, such as guilt and shame, exhibit marked postural differences [18] which might be observed in how you stand, walk, gesture, or otherwise behave.

Affective computers can begin to recognize these forms of emotional expression, despite the fact that there is still no solid definition of emotions, no resolution of “basic” or “discrete” emotions, and no agreed upon evidence for universal patterning. I will discuss these issues further below, but first let us consider some of the other complicating factors which exist.

#### 2.3.1 Complicating conditions

Studies attempting to associate bodily response with emotional state are complicated by a number of factors. For example, claims that people can experience emotions cognitively (such as love), without a corresponding physiological response (such as increased heart rate) are complicated by issues such as the intensity of the emotion, the type of love, how the state was supposedly induced (watching a film, imagining a situation) and how the person was or was not encouraged to express or suppress the emotion. The latter can be particularly subtle, as there are social “display rules” hidden in each culture, indicating when and where it is appropriate to express certain emotions.

Another complicating factor is that there may be physiological responses similar to those in an emotional state, but not corresponding to an emotion, e.g. heart rate also increases when exercising. However, a wearable computer that is trying to measure emotions might also have a camera and other sensors attached, e.g., step rate, so that it can recognize you are exercising. These basic activities can be built into the affective recognition model as conditioning variables. For example, in our current research with a wearable wireless affective head-mounted camera (for augmenting visual memory), we have found it more relevant to associate frame-rate not with just heart-rate, but with a function that combines heart-rate and step-rate.

Leidelmeijer overviews several conflicting studies in [27], reminding us that a specific situation is not equally emotional for all people and an individual will not be equally emotional in

all situations. Such studies point to hard-to-measure biochemical and cognitive influences in emotion, the latter of which are addressed more carefully below.

### 2.3.2 Personal patterning instead of universal patterning

One of the outstanding problems in trying to associate emotions with physiological responses is that sometimes different individuals exhibit different physiological responses to the same emotional state. Leidelmeijer [27] discusses the evidence both for and against universal autonomic patterning. The difficulties in finding consistent universal patterning mechanisms appear to make the outlook grim for constructing computers that can recognize affect.

However, my experience with signal-processing recognition systems leads me to suggest that this situation parallels that of another problem, the problem of constructing “speaker-independent” speech recognition systems, and therefore has a creative solution.

The goal of speaker-independent systems is to recognize what was said regardless of who said it. Even among people who use the same language, this goal is complicated by the fact that two people saying the same sentence produce different sound signals. The computer has difficulty separating the language part of the signal from the part of the signal that identifies the speaker and his or her expression. Consequently, the computer has a hard time recognizing what was said unless it was trained on the individual speaker, or on someone that sounds like that speaker.

Although it would be a terrific accomplishment to solve this universal recognition problem, it is unnecessary in affective computing. Negroponte pointed out years ago that an alternative solution is to solve the problem in a speaker-dependent way, so that your personal computer can understand you and your language; thereafter, your computer can translate what you said to the rest of the world.

Similarly, I believe that experiments in recognizing emotional state from observations of physical expression only need to demonstrate consistent patterning for an individual in a given perceivable context. It is only necessary that your personal computer be able to recognize your affect; it can then translate this information if you permit it to be communicated to others.

For example, the affect recognizer for one context may find it most reliable to use a blend of blood-flow and facial expression for one person, and a measure of skin-temperature, galvanic skin response, and respiration for another person. This adaptability greatly simplifies the problem, as a computer learning algorithm can be used to determine which features are the best predictors for each individual. Moreover, typically in adaptive scenarios a relatively small number of categories emerges, i.e., certain individuals tend to exhibit similar physiological responses, simplifying the amount of adapting the system has to do.

The individual’s personal computer will respond best if it is also able to perceive context (e.g., sense if you’re climbing stairs, if the room temperature changed, if you just read a news story about a tragic bombing, etc.) The computer can therefore identify autonomic responses conditioned on perceivable factors. For best performance, perceivable context should ultimately include not only the public milieu such as the comfort index of the weather, but also the private milieu – for example, the information that you have family in the town where the giant bomb just exploded. The priorities of your personal affective software agent need to shift with your affective state. For example, predicting your concern for your family as a top

priority, it might quickly alert you to more news on the bombing, or fetch phone numbers to help you get through to your family members wherever they are.

### 2.3.3 Studies out of the laboratory

The complications noted above have particularly plagued laboratory studies. For example, certain subjects might feel inhibited about expressing disgust during a laboratory study. Other subjects might find the situations in the study contrived, and exhibit a much smaller repertoire of emotions than they would experience in their natural world. Moreover, the “natural world” itself will differ among subjects – actors and musicians tend to exhibit emotions more readily than scientists and engineers.

As mentioned above, emotion theorists have tended to look for universal patterns, instead of for consistent patterning within an individual, conditioned on a context. With affective personal computers, understanding the individual response is most important first; after reliable data has been gathered for individuals, then comparisons can be made across populations.

Because most studies on emotion and cognition have been confined to artificial lab scenarios, they have been severely limited. Affective computers, particularly if wearable, would be able to measure emotional responses via sentic modulation wherever and whenever they occur, for both individuals and larger groups. Affective computing allows the laboratory to visit the subject, instead of the other way around.

## 2.4 Cognitive aspects of emotion

Humans are frequently conscious of their emotions, and we know from experience and laboratory study that cognitive assessment can precede the generation of emotions. Consequently, some prominent scientists have argued that cognitive appraisal is a *necessary* precondition for emotion. Although it is hard to “prove” that any human experience exists independent of cognitive events, there seems to be ample evidence that emotions can occur without prior cognitive appraisal [4], [7], [28]. In particular, the recent neurological evidence seems to support that emotions can “hijack” the cognitive centers of the brain [12]. Additionally, noncognitive biochemical events can strongly influence mood [4]. Only recently have scientists begun to unlock the secrets of hormonal chemistry, the role of neurotransmitters in depression, and other significant noncognitive contributors to human emotion.

It seems safe to conclude that both cognitive and physiological events can contribute to emotion, and vice-versa. Note there is no hard distinction between cognitive and noncognitive just as there is no hard line between the brain and the body; after all, thinking is both a cognitive and a physiological event; the mind-body separation is one of convenience, not physical reality.

A helpful distinction for sorting the “noncognitively-generated” and “cognitively-generated” emotions is made by Damasio [7] who distinguishes between “primary” and “secondary” emotions.<sup>10</sup> Damasio’s idea, which is also supported in much of the emotion theory literature, is that there are certain features of stimuli in the world that we respond to emotionally first, and which activate a corresponding set of emotions (and cognitive state) secondarily. Such emotions (such as startle upon hearing a loud bang, or the fear that causes an infant to retreat when a large object looms rapidly near) are “primary” and reside in the limbic system (most likely, the amygdala).

<sup>10</sup>Damasio’s distinction compares to that of “preattentive” and “postattentive” processing in the vision literature.

These are the innate emotions, Jamesian in their accompanying physical response. But these are not all. Damasio defines “secondary” emotions as those that arise later in an individual’s development when systematic connections are identified between primary emotions and categories of objects and situations. An example is grief, where physical responses occur in conjunction with cognitive understanding of an event (such as death of a loved one). For secondary emotions, the limbic structures are not sufficient; prefrontal and somatosensory cortices are also involved.

Damasio’s patients were unusual in that they had primary emotions but not secondary emotions. The “hard-wired” emotions were there, but the ability to generate appropriate cognitive emotions was not, even though the patients “knew” the appropriate emotional response to be generated.

The cortical involvement in the secondary emotions presumably helps construct important links connecting cognitively-recognized events to primary emotional responses. These links allow us to generalize primary events, such as flight from big looming objects, into cognitive constructs such as “stay out of the paths of big objects, such as cars, trucks and trains.” The cortical-limbic links set up during construction of secondary emotions can also be used to cognitively call forth a corresponding primary emotion at will.

The complex cortical activities available to humans probably also account for their ability to construct “self-conscious” cognitive emotions such as shame and guilt, which are not present in infants, but develop later in life [29]. Babies demonstrate a less complicated repertoire of emotions than cogitating adults, despite the fact that babies have not learned the social rules of suppressing emotions.

#### 2.4.1 Complicating conditions

A number of factors confound “purely cognitive” attempts to understand emotion. Several of the factors mentioned in the previous section are relevant to lab studies of cognitive emotions, especially social display rules and biochemical influences. In cognitive studies of emotion, subjects are usually asked to verbalize their emotional state, as opposed to (or in addition to) its being physically measured. However, the problem of attaching adjectives to emotions is immense [30].

Wierzbicka [31] has made one of the most comprehensive attempts to define emotion concepts in terms of universal semantic primitives such as “good” “bad” and “want,” resulting in a distinct script for each emotion concept in terms of a set of primitives. These scripts involve goals, and their achievement or lack of achievement.

Affective computers could simulate many of the competing goal-motivation theories of emotion by encoding them into scripts. Computers could then be put into social situations, like was done in the Stanford studies mentioned earlier. Affective computing could therefore provide a test bed for important cognitive-generative theories of emotion. However, affective computers cannot currently expect to measure cognitive influences; these depend on self-reports which are likely to be highly variable, and no one can read your mind (yet).

Under controlled environments, or with the assistance of some wearable acoustic and visual scene analysis, we may, however be able to measure the perceivable cognitive milieu, e.g., the room is tiny and hot, a stranger enters walking towards the subject rapidly, etc. We can also measure physiological responses such as facial expression, breathing, etc., with the more modes initially, the better. We should at least be able to measure physiologically those emotions which are already manifest to others. Both the cognitive milieu (including per-

ceivable environment) and physiological response are important in beginning to recognize and understand human emotions in natural situations.

As data is collected from a variety of situations, patterns may be found which would improve a computer’s ability to predict cognitive emotional responses to situations.

The successful predictive mechanisms, as in the physical case, is likely to depend on an individual. For example, someone who has never seen a neighborhood cat run over by a car might be deeply disturbed the first time they are close to such an event; however, the man who routinely removes animal carcasses from the roads may have no emotional response. The person’s history, values, and general emotional maturity combine to influence their cognitive responses.

Therefore, universal models of emotion-producing mechanisms are likely to be over-simplified and in need of customization when it comes to predicting individual responses. An affective computer could collect and analyze events and responses for an individual, and compare these pairings to a set of predictive models, determining which models give the most accurate predictions under different circumstances. Such adaptive data-collection abilities will be necessary to develop comprehensive theories of the cognitive aspects of emotion.

#### 2.4.2 A challenge in understanding creativity

An area for fruitful investigation is understanding how cognitive-affective links influence memory retrieval, and in conjunction, creativity. Improvement in memory retrieval has been found to occur when a person is in the same emotional state associated with the memory being retrieved [32].<sup>11</sup> Memory retrieval is a key activity in making free associations, which are important for analogical thinking and creativity. Hence, it is natural to expect to find relations between emotional states and creativity.

Damasio’s findings linked cortical constructs to primary emotions; consequently, we might expect cortical constructs for creative thinking and memory retrieval to also develop emotion links. The mechanism Damasio describes may therefore account for a separate idea, recently proposed by the computer scientist Gelernter [33]. Gelernter has suggested a phenomenon he calls “affect linking” which might play an important role in creativity. However, Gelernter suggested that this phenomenon arises primarily during what he termed “low focus” thinking, and not during what he termed “high focus” reasoning. However, as I have indicated, the neurological evidence indicates that emotion does play an important role in higher-level decision making; it is therefore not restricted to “low-focus” thinking as Gelernter muses.

### 2.5 A note on inducement of emotion

Certain physical acts are peculiarly effective, especially the facial expressions involved in social communication; they affect the sender as much as the recipient. – Marvin Minsky [34]

There is emotional inducement ever at work around us – a good marketing professional, playwright, actor, or politician knows the importance of appealing to your emotions. Aristotle devoted much of his teachings on rhetoric to instructing speakers how to arouse the right emotions in their audience [10].

Although inducement of emotions may be deliberate, it seems we, the receiver, often enjoy its effects. Certainly, we enjoy selecting a stimulus such as music that will affect our mood in

<sup>11</sup> A discussion of numerous affect-memory experiments, as well as some controversy surrounding them, appears in [1].



a particular way. We tend to believe that we are also free to choose our response to the stimulus. An open, and somewhat ominous question is, are we always free to do so? In other words, can some part of our nervous system be externally<sup>12</sup> activated to force experience of an emotion?

A number of theorists have postulated that sensory feedback from muscle movements (such as facial) is sufficient to induce a corresponding emotion. For example, Laird [32] divides people into “cueing” categories based on whether or not posturing themselves in a particular expression induces the corresponding emotional experience. Izard overviews some of the evidence for and against various sensorimotor claims [4].

Whether or not such sensorimotor inputs can induce emotion, they appear to at least be effective in maintaining and expressing emotion. Posture is correlated with expressions of self-esteem [4], [18]. A successful school of acting (after Michael Chekhov, student of Stanislavsky) is based on imagining emotive scenarios, and adjusting ones body position in accord with that emotion [36], [37].

Actors who excel at this method strengthen their association with their character’s emotional state. When the body’s emotional expression, e.g. an angry face and posture, agrees with the cognitive emotion, “my character is now angry,” the combined emotional experience is enhanced, “I feel angry.” Consequently, its communication to the audience is more powerful. Of course, these actors adjust their posture in accord with an initially cognitive goal. Hence, this is an example where emotions are initially cognitively-generated and the body-mind reinforcement intensifies and regulates the experience.

The body-mind reinforcement may also provide a subliminal way to induce emotion, perhaps by inducing involuntary eye saccades. Although the potential of such methods to induce emotion is unknown, the answers to questions like this may hinge on only a slight willingness<sup>13</sup> to be open to inducement.

The possibility of subliminal inducement may evoke disturbing thoughts of potentially harmful mind and mood control; or potentially beneficial mental enhancement and increased affective freedom. It is not an area to be entered into without considering both negative and positive aspects of how such new understanding could be used. As computers develop the ability to recognize affect, they potentially may be used for monitoring it, for both helpful and harmful purposes.

To what extent do neural, sensorimotor, motivational, and cognitive systems interact in emotion inducement and suppression? These are open research areas, and important for understanding both how natural helpful emotions are induced, and how harmful inducement can be minimized.

### 3 Affective Pattern Recognition

Although with the aid of new measuring devices we can distinguish many new activity levels and regions in the brain, we cannot, at present, directly access another’s thoughts or emotions. Instead, thoughts and emotions are communicated through words, gesture, music, and other forms of expression. The couplings between emotions and these forms are not well understood. What signals to measure, how to process these signals, and how to interpret them are open questions. These and many other hurdles must be overcome to give computers the ability to recognize affective states.

<sup>12</sup>External, in contrast with direct stimulation of the brain which is known to elicit various emotions [35].

<sup>13</sup>Perhaps this willingness may also be induced, ad infinitum.

However, computer recognition of affective states appears doable in many cases, via the measurement of sentic modulation. Note that I am not proposing one could measure affective state directly, but rather measure observable functions of such states. These measurements are most likely to lead to successful recognition when subjects do not deliberately suppress emotional expression. If one can observe reliable functions of hidden states, then these observations may be used to infer the states themselves. Thus, I may speak of “recognizing emotions” but this should be interpreted as “measuring observations of motor system behavior that correspond with high probability to one or more underlying emotions.”

Despite its immense difficulty, recognition of expressed emotional states appears to be much easier than recognition of thoughts. In pattern recognition, the difficulty of the problem usually increases with the number of possibilities. The number of possible thoughts you could have right now is limitless, nor are thoughts easily categorized into a smaller set of possibilities. Thought recognition, even with increasingly sophisticated brain imaging techniques, might well be the largest “inverse problem” in our world. In contrast, for emotion recognition, a relatively small number of simplifying categories for emotions have been commonly proposed.

#### 3.1 Basic or prototype emotions: key issues

##### 3.1.1 Categories or continuum?

Diverse writers have proposed that there are from two to twenty basic or prototype emotions. (See for example, [38], p. 8, [27], p. 10). The most common four appearing on these lists are: fear, anger, sadness, and joy. Plutchik [38] distinguished among eight basic emotions: fear, anger, sorrow, joy, disgust, acceptance, anticipation, and surprise. Ortony *et al.* provide a helpful summary of lists of basic emotions in their book [39]. Sometimes these “basic” emotions are defined to be essentially innate like Damasio’s primary emotions, but there is no consensus on their definition.

The actual existence of basic emotional states is disputed by some authors. Leidelmeijer [27] and Stein and Oatley [40] bring together evidence for and against the existence of basic emotions, especially universally, although I distinguish universality as a separate issue, addressed below.

Some authors have been less concerned with eight or so prototype emotions and instead refer primarily to continuous dimensions of emotion, such as negative or positive emotions. Three dimensions show up most commonly. Although the precise names vary, the two most common dimensions are “arousal” (calm/excited), and “valence” (negative/positive). The third dimension tends to be called “control” or “attention” addressing the internal or external source of the emotion, e.g., contempt or surprise.

The problem of not being able to precisely define categories is not restricted to emotions. It occurs all the time in cognition, pattern recognition, and so-called “fuzzy classification.” Nonetheless, I think the use of categories greatly simplifies recognition and communication of emotional state for both people and computers. A simplifying set of categories, chosen for their relevance to a particular practical domain, assists computers in beginning the difficult task of affect recognition. It is fitting that babies appear to have a smaller repertoire of emotions than adults.

The development of affective computing does not require resolution of these theoretical issues, although it should aid in their investigation. In a typical application or context, such as the piano tutor above, recognition of only a small set of emotions

can be of great benefit. Recognition of elements in the set does not imply that they are discrete or basic. The lack of consensus about the existence of precise categories of basic emotions does not interfere with the ideas I present below.

For affective computing, the recognition and modeling problems are simplified by either the assumption of a small set of discrete emotions, or the assumption of a small number of dimensions. Those who prefer to think of emotions as continuous can consider the discrete categories to be regions in a continuous space, or can adopt one of the dimensional frameworks. In a sense the choice of discrete or continuous states is like the choice of particles or waves in describing light. Ultimately, the one which is best depends on what you are trying to explain. Either representation, continuous or discrete, is commonly used in pattern recognition and is accompanied by helpful analysis tools.

### 3.1.2 Universality?

Much of emotion theory has been stymied on the issue of universality regarding the existence of emotion categories. In other words, if a set of emotions is truly “basic,” then shouldn’t they exist in all humans?

Like many questions in emotion theory, the study of this question is complicated by factors such as how the emotion was provoked, expressed, communicated, and labeled – after all, different languages may or may not use the same words for describing emotive phenomena. One of the potential benefits of affective computing lies in its ability to make measurements, conditioned on individuals, and on circumstances affecting them. Analytic tools can subsequently be used to search for universal patterns in the data.

### 3.1.3 Pure or mixed?

The debate in the literature about the purity of emotional states is another debate where experiments could be conducted with affective computers.

For example, Clynes’s exclusivity principle of emotional states [23] suggests that we cannot *express* one emotion when we are feeling another, e.g., we cannot express anger when we are feeling hope. Clynes emphasized the “purity” of the basic emotional states, and suggested that all other emotional states are derived from this small set of pure states, e.g., melancholy is a mixture of love and sadness.

Plutchik said that one can account for any emotion by a mixture of the principal emotions [38], and that emotions are rarely perceived in a pure state. This idea was captured by cartoonist Johnny Hart (and reprinted in [39]) in his “B.C.” cartoon illustrating an example of a mixed emotion: “seeing your long-lost dog come bounding up your freshly poured front sidewalk.”

The distinctions between the views of Plutchik and Clynes appear to be a matter of intensity and deliberate expression. If you are deeply involved in playing a mournful piece of music, you may be in a pure state of sadness. However, if you are thinking about lunch while playing the same piece of music, the measured state will likely not be as pure.

How does sentic modulation change as a person suppresses one strongly-felt state and tries to feel another? Could sentic measurements help people in identifying an emotion they are masking, such as in the expression of anger to hide fear? These are questions that affective computing could address, by constructing models of affective states and giving the computer the ability to recognize and record observations that correlate with these states.

An interesting set of tests could involve actors. Imagine a highly-skilled actor trying to play the role of a hopeful character when he is feeling anger right before the show. In order to deliberately express hope, he suppresses his anger, or overrides it with hope. If he is successful onstage in communicating hope, has he merely “forgotten” his anger, so that it will return after hope has finished its reign on stage, or is there a therapeutic effect that takes place during this performance? Measurements of his emotion before, during, and after the performance could be studied both for understanding the purity of emotions as well as for understanding their therapeutic effects.

Although “forgetting” sounds like a cognitive act, it has to occur bodily for emotions or the audience will sense conflict in the actor, instead of hope, and think him to be a bad actor. The intensity of the affective communication is a function of the actor’s sentic modulation – voice, face, posture, and more. As the actor deliberately brings all these modes into a consistent expression, not only is his communication more effective, but the theory is that he moves himself closer to a pure state of emotion. If Clynes’s theory holds, then the purer the emotional state, the more effective and affective the actor’s communication will be.

Theories such as this, examining the purity of emotions and of their influences, could be tested empirically with an affective computer that recognizes emotional states, as described next.

## 3.2 Modeling affect

How does the computer represent emotions? Obviously, current computers do not have the equivalent of a limbic brain and a cortical brain, or the biochemical washes that connect these regions, or the pyramidal and non-pyramidal systems, and so forth. Computers have “bodies” but they are currently not affective.

In this section I will suggest both computational and rule-based models for representation of emotions. This section contains technicalities which can be skipped by those who are not engaged in designing systems that perform emotion analysis (recognition) or emotion synthesis (prediction, and generation).

The models will roughly be divided into three types: computational models for discrete states, computational models for emotion spaces, and rule-based models.

### 3.2.1 Discrete affective state models

This section considers possible models for the discrete, hidden paradigm of emotional states. Figure 1 shows an example of one possible model, the Hidden Markov Model (HMM). This example shows only three states for ease of illustration, but it is straightforward to include more states, such as a state of “no emotion.” The basic idea is that you will be in one state at any instant, and can transition between states with certain probabilities. For example, one might expect the probability of moving from an interest state to a joy state to be higher than the probability of moving from a distress state to a joy state. The actor described in the previous section might circulate among a hope state, an angry state, and a “mixed” state.

The HMM is trained on observations, which could be any measurements of sentic modulation varying with the underlying states, such as changes in voice inflection or heart-rate. The input at any time is these observations; the output can be either the state that the person is most likely in, or it can be identification of an entire HMM configuration, thereby recognizing a pattern of emotional behavior.

In the latter case, there would be a family of configurations, one corresponding to each useful emotional behavior, or each individual’s characteristics for a given behavior. In either case,

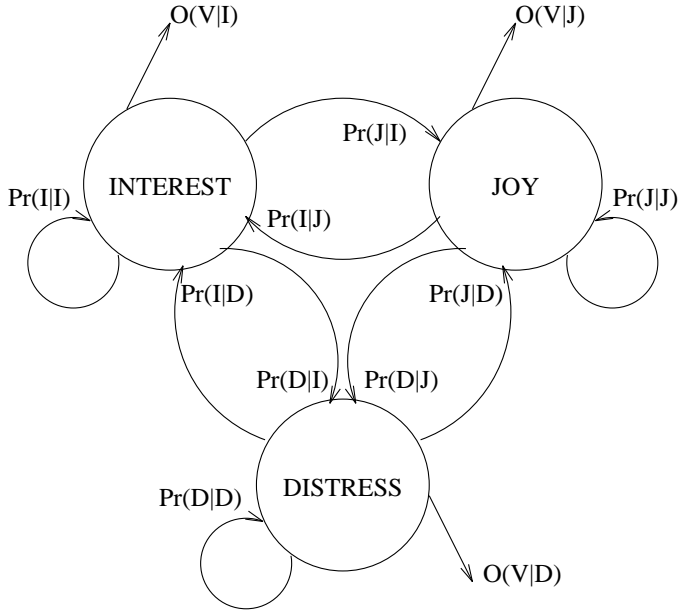


Figure 1: The state (e.g. Interest, Distress, or Joy) of a person cannot be observed directly, but observations which depend on a state can be made. The Hidden Markov Model shown here characterizes probabilities of transitions among three “hidden” states, (I,D,J), as well as probabilities of observations (measurable essential forms, such as features of voice inflection, V) given a state. Given a series of observations over time, an algorithm such as Viterbi’s [41] can be used to decide which sequence of states best explains the observations.

the HMM states do not have to correspond to pure emotional states as illustrated in Fig. 1, but may correspond to even more fundamental building blocks, perhaps identified by the computer as it works to fit the data.

In either case, different HMM’s can be trained as a function of environmental or social context, e.g., driving a car in the country vs. in the city at rush hour, going out with an old friend vs. on a blind date. There could also be different HMM’s based on timing relative to a hormone cycle or to exam season. Hence, the probabilities, states, and their connections may vary depending on a variety of factors, ultimately determined by the intended use of the affective state recognizer. The model is flexible, and can accommodate many variables of interest in emotion theory.

HMM’s can also be adapted to represent complex mixture emotions. One could design experiments to see which mixture combinations made the best predictors of individual emotional responses. For example, an HMM of interest in stressful learning situations might represent a causal sequence of attention/high-arousal, followed by distress and then by relief states. The choice of states could be made by clustering physiological variables, and assigning each cluster its own state.

Static mixtures may also be modeled (and tailored to an individual, and to their context) by explicit “mixture models” such as the cluster-based probability model of Popat and Picard [42]. In such a case, high-dimensional probability distributions are learned for emotional states or their mixtures based on the values of or functions of the values of the physiological variables.

The input would be a set of observations, the output a set of probabilities for each possible emotional state. Such a probabilistic formulation can also provide “fuzzy” classifications of emotional state, where someone’s state might be described as 75% joy and 25% distress, as perhaps in the B.C. cartoon mentioned above.

The models I’ve described for affect recognition can be used not only to represent emotional states and behaviors, but also to predict and “synthesize” them. The prediction process is one of partial recognition: first, determine which model (and parameters) best fits your current observations, and then see what state that system would most likely synthesize next. Such a model-based prediction would give a “likely” outcome, but could not predict with 100 % certainty what would actually happen.

Notice that emotional state synthesis by the computer involves no sentic modulation, but only the generation, in name, of a state or behavior. We can synthesize samples from a probabilistic mixture model to obtain “typical” behaviors according to that model, e.g., when these nine variables lie in this range then the prediction is that grief will arise, but we cannot make a computer cry or laugh yet (although they are getting better at voice inflection, and it is a small step to use a synthesized model state to drive a vocal output.) Synthesis of emotion is a topic I will revisit in Sect. 3.2.4.

Numerous other tools from pattern recognition are also likely to be useful in affect recognition. Artificial neural nets can perform a variety of recognition tasks and can function like mixture models; hence, they should be useful for emotional state modeling. Neural nets and related models such as the M-Lattice [43] can also model certain nonlinear dynamical systems. Camras [44] has proposed that dynamical systems theory be considered for explaining some of the variable physiological responses observed during basic emotions, but has not suggested any models. Freeman has modeled olfaction with dynamical systems and proposes the importance of this approach for modeling limbic influences on intention and motivation in his book *Societies of Brains* [45], but he has not proposed any computational models for the latter.

### 3.2.2 Continuous affect models and “eigenmoods”

Instead of assuming discrete states, sometimes it is more appropriate to start with the data and perform factor analysis or an eigenvector decomposition to discover its underlying dimensions. In this case, one tries not to uncover discrete states, but rather continuous axes which describe the most important variations in the data.

For example, in the case of an eigenvector decomposition, we would start with a variety of observations of sentic modulation measurements, use these to form a covariance matrix, and then find the eigenvectors which diagonalize this covariance. The eigenvectors corresponding to the top eigenvalues could be kept, resulting in a space of an arbitrarily small number of dimensions. Discriminants could be computed to determine how well distances in the resulting eigenspace corresponded to perceived differences in the corresponding emotional states.

The most useful dimensions in a person’s emotion eigenspace might correspond to “eigenmoods,” providing building blocks for emotional expressions. As new expressions are observed and projected into this space they are described in terms of these underlying eigenmoods. The eigenmoods might correspond to either either pure or mixture emotions; the eigenspace model works with either interpretation.

Note that in modeling, any signal can be decomposed into basis components; therefore, one can always find sub-components

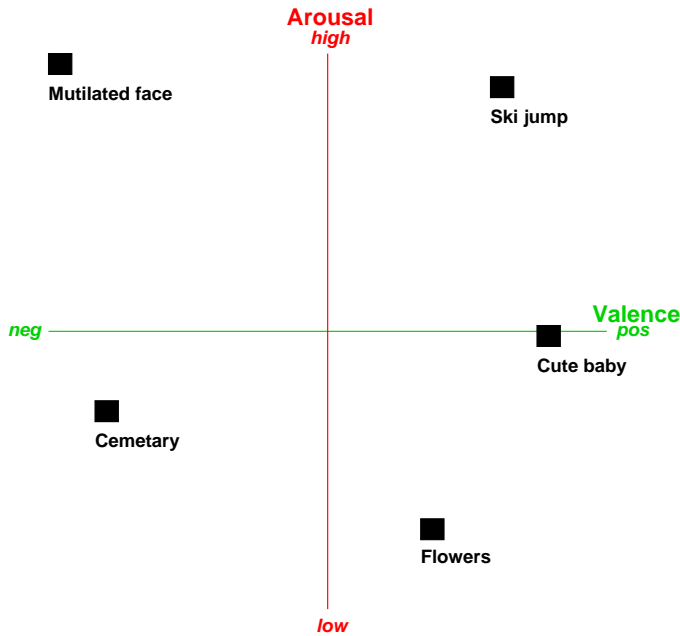


Figure 2: Two dominant dimensions of emotion, here used to locate the affective response of certain photo contents, as per the studies of Lang [46].

– a so-called mixture, even if the signal is “purely” of one kind. Consequently, from a modeling perspective, the theoretical issue of “pure” vs. “impure” emotional states is not problematic.

Eigenspaces constructed from measurements of sentic modulation over the course of a person’s day-to-day activities should be compared to the spaces found by factor analysis in emotion theory. The spaces could be estimated under a variety of conditions, to better characterize features of emotion expression and their dependencies on external (e.g., environmental) and cognitive (e.g., personal significance) factors. Trajectories can also be characterized in these spaces, to study the dynamic aspects of emotion.

An example of the most common dimensioned space for emotions is shown in Fig. 2. Such a space might be more useful than an HMM if the affective computer is trying to gauge how positive or negative users feel about a new software product, for example.

Given one of these dimension-space models trained on motor outputs corresponding to known emotional states, then features of unknown motor outputs can be collected and used with signal processing tools such as maximum a posteriori decision-making to recognize or classify a new emotion. Because the recognition of emotional state can be set up as a pattern recognition problem, a variety of techniques are available for its solution [47], [48].

### 3.2.3 Cathexis in computing

Although most computer models for imitating mental activity do not explicitly consider the limbic response, a surprisingly large number implicitly consider it. Werbos [49] writes that his original inspiration for the backpropagation algorithm, extensively used in training artificial neural networks, came from trying to mathematically translate an idea of Freud.

Freud’s model began with the idea that human behavior is

governed by emotions, and people attach cathexis (emotional energy) to things Freud called “objects.” Quoting from Werbos [49]:

According to his [Freud’s] theory, people first of all learn cause-and-effect associations; for example, they may learn that “object” A is associated with “object” B at a later time. And his theory was that there is a backwards flow of emotional energy. If A causes B, and B has emotional energy, then some of this energy flows back to A. If A causes B to an extent W, then the backwards flow of emotional energy from B back to A will be proportional to the forwards rate. That really is backpropagation....If A causes B, then you have to find a way to credit A for B, directly. ...If you want to build a powerful system, you need a backwards flow.

The use of some form of feedback is a significant part of most computer learning methods today. Usually, the feedback is via a set of positive or negative examples provided by a user. A more powerful learning system might augment the positive and negative examples with affective responses from the user.

In other words, now the computer is trying to learn from the user. Suppose it tries to learn your preferences for art to hang in your livingroom. As you browse a database of images looking for suitable pictures, it detects your response to different images. After you have indicated your favorites to it, it could try to infer which art features (e.g., artist, color, texture, content) and which sentic features (e.g. facial expression, heartrate, skin response) were the best predictors of what you liked. It might even be able to associate certain categories of images with certain categories of your affective responses. In the future, it might watch for occurrences of those same features, to help you save time locating things you like, or locating things that suit your present mood.

Such an affective learning agent might be a valuable assistant to your interior decorator, or to the personalized environments of the future, that might favorably customize your hotel, conference, or retreat surroundings to win your business.

A variation on the HMM above could also be used to incorporate affective feedback. The recent Partially Observable Markov Decision Processes are set up to give a “reward” associated with executing a particular action in a given state [50], [51], [52]. These models permit observations at each state which are actions [53]; hence, they could incorporate not only autonomic measures, but also observations of your behavior.

### 3.2.4 Rule-based emotion models and their limitations

The focus in the previous section was affect recognition using mathematical models, which, in some cases, can also be used for emotion synthesis and prediction. Alternatively, non-mathematical rule-based models may be used.

Some work on rule-based models of emotion has been done in the AI community, where the emphasis has been on writing scripts to produce states that are labeled with various emotions (See Pfeifer [54] for a nice overview.) The AI emphasis has been on rule-based synthesis of cognitive states *in* the computer, which receive emotional names – for example, if the computer has the goal of getting you to fill out a form, and you do not fill it out after it repeatedly asks you to do so, then it might enter a state called “frustrated.”

Although I have not yet seen work with rule-based models for emotion recognition and prediction, this would be possible given a good synthesis model coupled with signal processing

and pattern recognition to convert the human signal into the language of the script (or vice versa.) The sentic modulation measurements would first undergo translation from the numerical to symbolic form, or better, the AI models would adapt to learn the sentic modulation language. In either case, it will still be important to use some of the techniques described above.

Therefore, the models of Pfeifer and other cognitive-motivational generators of emotion (such as the proposed scripts of Wierzbicka in [31] which could be programmed) potentially not only could run a script to generate a cognitive-emotional state, but also could identify which components of the script are satisfied by a set of observations of a human. In this way, the script could recognize someone else's emotional state or predict a forthcoming emotion.

Recognition (analysis) models in general do not have to be capable of synthesis. This holds true for modeling emotion, allowing emotion to be recognized by a computer without being synthesized. However, it is interesting to point out that even with emotion analysis and synthesis, a rule-based system is still likely to be defective.

Consider again the frontal-lobe patients of Damasio. These patients can recognize emotions, and can predict (synthesize) that they should feel certain emotions (such as horror) given certain events (such as seeing a picture of a head being chopped off). They have cognitive analysis and synthesis capabilities, but still cannot behave normally. What they are missing is the physiological response that must work in concert with cognitive evaluation.

Based on the findings of Damasio, it seems to be true that not only are emotion analysis and synthesis an essential part of a future intelligent computer, but two kinds of synthesis are needed – the counterparts of both the cognitive and the physiological systems. The cognitive component might be fit well by a set of laws or rule-based models. However, these laws would not be enough; relying only on them to simulate emotion would be a modern form of Descartes' error in separating body and mind. Rather, the computer must also have a second component – a “bodily component.” This component would not follow the logically neuronal-inspired mechanisms of present computers, but would aim to function as the biochemical flow that floods the brain, activating and focusing attention, determining salience, biasing decisions, and driving thoughts. This bodily component in the computer could both interact with, and if needed, override the laws and rules of the reasoning component – march to its own tune, so to speak, acting as the equivalent of our human “songs.”

If these two components are not present, then not only can we not expect to see creative computers, but we can expect them to remain decision-impaired. The analogy I have made is of rule-based machines to patients like those of Damasio, who are highly rational but unable to function normally in society. However, Damasio's patients at least have primary emotions; it is the secondary ones they are missing. Therefore, this “bodily” component I am proposing must do more than synthesize a (primary) emotional state for the computer, it must interact, as do secondary emotions in a human, with the rule-based reasoning of the machine. The rule-based decisions and emotional biases must operate in concert.

To summarize, the emotion component in the computer should not be a mere script for generating states, or mere random perturbations to bias an otherwise rule-based system. Instead, the emotion component must be closely coupled into the machine's intelligence, able to learn, and able to “run the show” from time to time.

### 3.3 “What do I appear to be feeling? What am I feeling? Where am I?”

Let's pause and consider another scenario, involving human affect analysis and synthesis.

Suppose you hear the winning lottery number, and recall that it is your friend's favorite number and she has played it every week for the last ten years. You cognitively predict (rule-based synthesis) that when she walks in, she will be elated. Your own physiological system, in anticipation, might synthesize a correspondingly excited emotional state.

If she walks in looking terribly distraught, and saying she heard the winning number, then you feel surprise, recognizing that your prediction is wrong. Then you (cognitively) learn that she cannot find her ticket. After you talk to her and learn that her ticket blew away in the wind, then you may synthesize a new emotion, such as compassion.

The conflict which occurred between the synthesis model's prediction and the analysis model's recognition is not only acceptable, but also flags an *interesting* event, one which is likely to generate a state of higher arousal, perhaps priming the physiological system for a faster response in case of a life-threatening emergency or life-enhancing opportunity.

A computer could also potentially flag these events, although a computer has apparently never been given a fast-responding physiological system for the purpose of avoiding life-threatening emergencies or to pursue life-enhancing opportunities. The physical correlates in a computer would likely take the form of a separate, parallel emotion synthesis model, capable of generating interrupts of priority sufficient to override or “hijack” the cognitive system if needed.

Coupling between synthesis and analysis models has some imminent practical applications. For example, in an existing environment such as ALIVE [55], a synthesis model could adjust each software agent's posture, facial expression, and gestures to reflect its (synthesized) emotional state. The state itself could also be re-synthesized (via the cognitive scripts mentioned earlier) as the agent recognizes the states of humans interacting with it in this virtual world.

In such an interactive system, a human might “role-play” different methods of emotional interaction. Hence the system becomes a test-bed for new strategies or games involving affective communication. More importantly, perhaps, such a test-bed provides a safe and controllable environment for exploring the nature and development of emotional intelligence, which, according to Goleman, can be learned [12].

Another important example of coupling synthesis and analysis is in the case of a speaking-impaired human relying on a speech synthesizer. Such people are usually limited to one inflection-less digital voice. Control over affect in synthetic speech is particularly important for these people [56]. With an affective wearable computer that senses sentic modulation, and couples it to a synthesis model, parameters for voice inflection could be synthesized and fed *directly* into their speech synthesizer. The result would allow a speaking-impaired individual to, for example, express anger with her voice by merely feeling angry. The coupling could allow non-verbal expression to drive the synthetic verbal expression.

If sentic modulation synthesis is coupled to the “inverse problem” of sentic modulation analysis (recognition), then in the ultimate modeling scenario, the emotion synthesis model could synthesize the affective responses of a body separated from its brain, as in Dennett's fantastic story “Where am I” [57]. Likewise, to the extent that sensorimotor stimuli affect and reinforce the cognitive state of emotion, a body's sentic modulation could

(after analysis) drive the brain into a corresponding cognitive state. Although this scenario is unthinkable with real humans and impossible with foreseeable technology, it could potentially be simulated with the aid of affective computers. Such admittedly far-out scenarios are nonetheless intriguing for the investigation of better theories about cognition, consciousness, and mind-body interactions.

## 4 Things Better Left Unexplored?

I'm wondering if you might be having some second thoughts about the mission – Hal, in the movie *2001: A Space Odyssey*, by Stanley Kubrick and Arthur C. Clarke

A curious student posed to me the important, infrequently-asked question, if affective computing is not a topic “better left unexplored by humankind.” At the time, the worst potential danger in my mind was that computers might be used to maliciously induce emotions. My inclination was to respond, “Emotion manipulation for both good and bad purposes is already commonplace, look at music, cinema, marketing, and politics. Wouldn't affective computers at least help educate people so that they can be more in control?”

However, questions like this do not have such easy answers; it is over-simplifying to say that improving people's understanding is the best solution. Although the Marie Curie quote which leads off this paper most succinctly summarizes where I lean, it would be irresponsible to not consider the potential harm as well as good that can come from this new area of research. Affective computing has potential for both good use and for misuse. Affective computers convert personal emotional expression into bits, making this information subject to all the problems associated with digital media, including issues of broadcast, privacy, and even copyright.

### 4.1 A dilemma

There are many beneficial reasons to pursue the development of affective computers. Nonetheless, their development raises the following dilemma:

*Can we create computers that will recognize and express affect, exhibit creativity, intelligent problem solving, and empathy, and never bring about harm by emotional actions?*

To elaborate, I will briefly present two scenarios from artists, viz. fiction writers and movie producers, who have presaged scientists in considering what may happen when computers receive emotions. Subsequently, I will discuss more imminent concerns.

The first scenario comes from Asimov's “The Bicentennial Man” [58]. Asimov subjects his affective robots to three laws of behavior to prevent them from bringing harm to people. His laws put human life above the self-preservation of the robot. However, his laws are not infallible – one can propose logical conflicts where the robot will not be able to reach a rational decision based on the laws. Indeed, his robots could be rendered completely ineffectual by getting into states where they could not reach a decision due to rule-based conflicts. Without a separate emotion system that can determine saliency, and ultimately override rules, a law-based robot is severely handicapped in its decision-making ability, not too unlike Damasio's frontal-lobe patients.

A somewhat more sinister scenario of an emotional machine occurs in the science fiction classic “2001: A Space Odyssey.”<sup>14</sup> A HAL 9000 computer, “born” January 12, 1997 (in the novel)

<sup>14</sup>The film was based on the 1965 screenplay by Kubrick and Clarke; Clarke's novel came out afterward in 1968 [59].

is the brain and central nervous system of the spaceship Discovery. The computer, who prefers to be called “Hal,” has perceptual abilities which emulate those of a human. Hal is a true “thinking machine,” in the sense of mimicking both cognitive and emotional functions. Humans who interact with Hal recognize his emotional abilities, as evinced in this exchange between a reporter and crewman of the Discovery:

Reporter: “One gets the sense that he [Hal] is capable of emotional responses. When I asked him about his abilities I sensed a sort of pride...”

Crewman: “Well he acts like he has genuine emotions. Of course he's programmed that way to make it easier for us to talk with him. But whether or not he has real feelings is something I do not think anyone can truly answer.”

As the movie unfolds, it becomes clear that the articulate Hal is capable of both expressing and perceiving emotion:

“I feel much better now.”

“Look, Dave, I can see you're really upset about this.”

But Hal goes beyond expressing and perceiving emotion. In the movie, Hal appears to *have* fear of being disconnected, as indicated not just by his spoken expression, but also by his reactive behavior. The novel indicates that Hal experiences internal conflict between truth and concealment of truth. This conflict results in Hal killing all but one of the crewmen.

Hal is more than a thinking and feeling machine; not only can he pass the Turing test, but also he can kill the person administering it.

The fictional message has been repeated in many forms and is serious: a computer that can express itself emotionally will some day act emotionally, and the consequences may be tragic.

Objection to development of “emotional computers,” based on fear of the consequences, parallels the “Heads in the Sand” objection, one of nine objections playfully proposed and refuted by Turing in [9] to the question “Can machines think?” But fear of the consequences is to be balanced against the practical benefits that should appear, given the importance of the limbic (emotional brain) role in thinking. Cytowic, talking about how the limbic system efficiently shares components such as attention, memory, and emotion, notes:

Its ability to determine valence and salience yields a more flexible and intelligent creature, one whose behavior is unpredictable and even creative. [6]

Today, with the ever-increasing information available to machines (and computer software agents), it is more important than ever for a computer to be given the ability to determine valence and salience. It is also commonly agreed that creativity and flexibility are necessary components of intelligence [60]. However, how to construct such qualities based on AI-style rules or without gratuitous randomness has so far eluded scientists. I think that the construction of such qualities will require mechanisms that duplicate both limbic abilities and cortical abilities.

I have argued a variety of reasons why intelligent machines will need emotions. The issue is no longer “is emotion necessary for intelligence,” nor even “how might machines analyze and synthesize emotions.” The evidence supports an answer of “yes” to the first question, and I have explained mechanisms for proceeding with the second. Instead, the issue at hand is do we want to give computers the ability to *have* emotions, the kind that can hijack their reasoning systems, and ultimately produce the unpredictable behavior that is the hallmark of creativity?

Can we build a machine and give up control over it? In other words, are we willing to give it free will to make value-based, emotional decisions? Such a machine would be guided by, but ultimately not constrained by, the ethics or mores which we give it. This machine would be free, for example, to deny our role in its creation.

## 4.2 Unemotional, but affective computers

Man’s greatest perfection is to act reasonably no less than to act freely; or rather, the two are one and the same, since he is the more free the less the use of his reason is troubled by the influence of passion. – Gottfried Wilhelm Von Leibniz [61]

Although expressing and recognizing affect are important for computer-human interaction, building emotion into the motivational behavior of the computer is a different issue. In fact the word “emotional” when it refers to people or to computers, usually connotes an undesirable reduction in rationality. Do you want to wait for your computer to *feel interested* before it will listen to you? Or, perhaps it would be beneficial if people who bombard you with email had to make the content sufficiently interesting before your computer would accept it?

Interestingly, in the popular series *Star Trek, The Next Generation*, the affable android “Data” was not originally given emotions, although he was given the ability to recognize them in others. Data’s evil android brother, “Lore,” had an emotion chip, and his daughter developed emotions, but was too immature to handle them. Although both Data and his brother had the ability to kill, Data evidently could not kill out of malice. A later episode focused on the maturity process needed for Data to deal with the replaced emotion chip. The process parallels what we expect with the development of emotional intelligence in humans [12].

One might argue that computers should not be given the ability to kill. But it is too late for this, as anyone who has flown in a commercial airplane acknowledges. Alternatively, perhaps computers with the power to kill should not have emotions,<sup>15</sup> or they should at least be subject to the equivalent of the psychological and physical tests which pilots and others in life-threatening jobs are subject to.

The fact remains, giving computers the ability to recognize, express, and “have” emotions is only the beginning of greater issues of how should they use these emotions. I will touch on this again below in Sect. 4.5.

A number of social and philosophical questions are raised by affective computers that have emotions: Should we follow the human paradigm and allow machines to express an emotion that is different from what they are experiencing internally? Should the affective computer have separate channels of expression, like we have in our pyramidal and non-pyramidal systems? Alternatively, should we give computers abilities humans do not have, such as the ability to have multiple pure emotions – or is this impossible without it having parallel self-awareness systems, multiple consciousnesses, multiple possible parallel personalities? What should the nature of the computer’s self-awareness be? Should we, the maker of the emotional machine, be the only one allowed to peer inside and see its true internal state? These are but a few of the issues raised

<sup>15</sup>Although I refer to a computer as “having emotions” I intend this only in a descriptive sense, e.g., labeling its state of having received too much conflicting information as “frustration.” I doubt electronic computers will *have feelings* as humans do, but I recognize the parallels in this statement to debates about machines having consciousness.

Computer	Cannot express affect	Can express affect
Cannot perceive affect	I.	II.
Can perceive affect	III.	IV.

Table 1: Four categories of affective computing, focusing on expression and recognition.

in affective computing; further discussion of these is outside the scope of this paper.

Clearly, computers could benefit from development of ethics, morals, and perhaps also of religion.<sup>16</sup> These developments are important even without the amplifier of affect. But computer scientists are not equipped to deal with these issues on their own, much less to create machines with these abilities. Scientists’ understanding of these great human issues has barely moved past that of Mary Shelley’s *Frankenstein*. Affective computing needs input from outside the cognitive and computer sciences.

## 4.3 Four cases for the present

The imaginary scenarios above take us far from present reality; as provocative as they may be, they remain science fiction. For the rest of this paper I would like to focus on more immediately available cases of affective computing.

Four such cases are summarized in Table 1. Of course these are not all possible cases; for example, I omitted the rows “Computer can/can not induce the user’s emotions” as it is clear that computers (and all media) already influence our emotions, the open questions are how deliberately, directly, and for what purpose? I also omitted the columns “Computer can/can not act based on emotions” for the reasons described above. The ethical and philosophical problems to be addressed before the development of such “emotionally-based computers” go beyond the scope of this paper; these possibilities are not included in Table 1 or addressed in the applications below.

This leaves the four cases described below:

- I. Most computers fall in this category, having less affect recognition and expression than a dog. Such computers are neither personal nor friendly.
- II. This category aims to develop computer voices with natural intonation, and computer faces (perhaps on agent interfaces) with natural expressions. When a disk is put in the Macintosh and its disk-face smiles, users may share its momentary pleasure. Of the three categories employing affect, this one is the most advanced technologically, although it is still in its infancy.
- III. This category enables a computer to perceive your affective state, which I have argued is important if it is to adjust its response in ways that might, for example, make it a better teacher and more useful assistant. This category allays the fears of those who are uneasy with the thought of emotional computers, in particular, if they do not see the difference between a computer expressing affect, and being driven by emotion.

<sup>16</sup>The latter raises particularly intriguing questions – e.g., Should they fear only their maker’s maker?

IV. This category maximizes the sentic communication between human and computer, potentially providing truly “personal” and “user-friendly” computing. It does *not* imply that the computer would be driven by its emotions.

#### 4.4 Affective symmetry

In crude videophone experiments we wired up at Bell Labs over a decade ago, my colleagues and I learned that people preferred seeing not just the person we were talking to, but also the image they were seeing of us. Indeed, this “symmetry” in being able to see at least a small image of what the other side is seeing is now standard in video teleconferencing.

It is my opinion that a similar form of symmetry should be considered for computers in categories III and IV (i.e., infer hidden emotional state based on physiological and behavioral observations). Your computer should let you see what it is reading as your state.

More precisely, affective interaction with a computer can give a person direct feedback that is usually absent in human interaction. The “hidden state” models proposed above can reveal their state to us, indicating what emotional state the computer has recognized. Of course this information can be ignored or turned off, but my guess is people will leave it on.

This feedback not only helps debug the development of these systems, but is also useful for someone who finds that people misunderstand his expression. Such an individual may never get enough precise feedback from people to know how to improve his communication skills; in contrast, his computer can provide ongoing personal feedback.

However, as computers receive the ability to synthesize internal emotional states, should these be observable by other computers? Or possibly observable only through some imperfect expressions, as is the case for humans? What are the implications for communication if emotion is unambiguous? It seems that imitating the human paradigm, where computers could mask their emotions from each other, would lead to interesting developments, for example, a software agent that was especially talented at negotiation.

Consider the issue of whether there should be an asymmetry, giving humans unequivocal access to read internal synthesized computer states. (Asymmetry, since I don’t think computers will achieve 100% recognition of human emotional states.) In 2001, if Hal’s emotional state were observable at all times by his crewmates, then they would have seen that he was afraid as soon as he learned of their plot to turn him off. If they had observed this fear and used their heads, then the tragic 2001 storyline would not work.

Instead, Hal illustrates the case of a computer that could hide its emotional state better than most people. Speaking in the spirit of its fictional character, it could “have its feelings hurt” and not let the humans know. I expect computers will indeed someday have the ability to feign emotions better than people. If so, a possible preventive step for disasters and miscommunication would be to prohibit the machine from hiding its emotions. Such a constraint, however, conflicts with the goal of giving it the full freedom required in true creativity.

#### 4.5 A distinction: affect recognition vs. intelligent response

Before proceeding with applications in the next section, there is an important distinction to be made regarding the scope of affective computing. Just about any application involving an affective computer will require attention to the following three issues:

1. What is the relevant set of emotions for this application?
2. How can these best be recognized/expressed/modeled?
3. What is an intelligent strategy for responding to or using them?

Affective computing primarily equips the computer to handle the first two issues. The third issue requires domain-specific knowledge, beyond affect analysis and synthesis, and into issues not just of emotional intelligence, but of combining emotional intelligence with the general goals of the application.

As one example, consider again the affective piano tutor scenario. An appropriate set of states for the computer tutor to recognize might include distress, interest, and pleasure. The computer might recognize these states from watching your face, gestures, posture, or measuring other responses of your autonomic nervous system. Affective computing in this application primarily enables pattern recognition of human emotional expression.

However, how should the computer respond when you make an error the first time? The Nth time? When you do something well? How should it adapt its responses to optimize your learning experience? Merely adapting to “always please” the user is naive, and conjures up the soma-dependent society of Huxley’s *Brave New World* [62]. Indeed, the answers to these questions go beyond affective computing, into questions of learning, epistemology, and more.

In some cases, this third issue opens up social and ethical questions. For example, suppose that an automobile’s affective steering wheel senses that you are unusually stressed or angry. (Many automobile accidents are caused by people who are angry or upset.) Recognition of your dangerous state is the job of affective computing; how the intelligent vehicle should respond is potentially a legal issue.

I will offer ideas below as to how affective computing might be used to help develop new theories to deal with this third issue. However, success with the third issue will require more than emotional intelligence.

## 5 Applications of Affective Computing

Let’s move past the philosophical issues above, and into examples of imminent practical applications. The rest of this paper considers potentially beneficial scenarios for applying affective computing. All of them assume the involvement of one or more human participants, who willingly participate in affective communication.

The examples below come mainly from cases II, III, and IV in Table 1 where computers can perceive and/or express affect. The scenarios assume modest success in relating observations of an individual’s sentic modulation with at least a few appropriate affective states, which is the most difficult new hurdle for affective computing research. However, the hurdle has already been leaped in some of the cases described below.

### 5.1 Entertainment

Why I do so well is I induce emotion. – Sean D. Tucker, American aviation artist [63]

One of the world’s most popular forms of entertainment is large sporting events – whether it is an outdoor air show, the Olympics, the World Series, Super Bowl, or any number of other large gatherings, there is an excitement in the air when fans come together to watch athletes perform.

One of the pleasures that people receive from these events (whether or not their side wins) is the opportunity to freely ex-



press intense emotions. I would not be surprised if the stereotyped “unemotional American male” might “need” these events psychologically; a stadium is one of the few places where such a person can yell and jump up and down, releasing emotion, without appearing childish. Emotional expression is acceptable for sports athletes and spectators. Indoors in front of a TV sports game, a mature adult can yell and jump with emotional expression that otherwise would be disdained. Emotions, and their expression, are a healthy aspect of entertainment.

Do you feel like I do? ... Do you *feel* like I do? –  
Peter Frampton

Although I’m not a fan of Peter Frampton’s music, I am still moved by the tremendous response of the crowd in his live recorded performance where he asks them this question repeatedly, with increasingly modified voice. Each time he poses the question, the crowd’s excitement grows. What causes such a tremendous emotional response from a crowd? Are they mindless fans who would respond the same to a mechanical repeating of the question, or to a rewording: “do you *think* like I do?” Or, is there something more fundamental in this crowd-arousal process?

I recently participated in a sequence of interactive games with a large audience (SIGGRAPH 94, Orlando), where we, without any centralized coordination, started playing Pong on a big screen by flipping (in front of a camera, pointed at us from behind) a popsicle stick that had a little green reflector on one side and a red reflector on the other. One color moved the Pong paddle “up,” the other “down,” and soon the audience was gleefully wagging their sticks to keep the ball going from side to side. Strangers grinned at each other and people had fun.

Pong is perhaps the simplest video game there is, and yet it was significantly more pleasurable than the more challenging “submarine steering adventure” that followed on the interactive agenda. Was it the rhythmic pace of Pong vs. the tedious driving of the sub that affected our engagement? After all, rhythmic iambs lift the hearts of Shakespeare readers. Was it the fast-paced unpredictability of the Pong ball (or Pong cat, or other character it changed into) vs. the predictable errors the submarine would make when we did not steer correctly? What makes one experience pleasurable more engaging than another?

Clynes’s “self-generating principle” indicates that the intensity of an emotional state is increased, within limits, by the repeated, arrhythmic generation of essentic form [23]. Clynes has carried this principle forward and developed a process of “sentic cycles” whereby people (in a controlled and voluntary manner) may experience a spectrum of emotions arising from within. The essence of the cycles is supposedly the same as that which allows music to affect our emotions, except that in music, the composer dictates the emotions to you. Clynes cites evidence with extensive numbers of subjects indicating that the experience of “sentic cycles” produces a variety of therapeutic effects.

Good entertainment may or may not be therapeutic, but it holds your attention. Attention may have a strong cognitive component, but it finds its home in the limbic system as mentioned earlier. Full attention that immerses and “pulls you in” becomes apparent in your face and posture. Affective computers might measure these responses to different forms of entertainment, providing researchers with signals that can be correlated with other measures of mental and physical health benefits. Just as studies involving test groups on meditation (a form of focused attention) have shown certain improved health

benefits, studies on other forms of attention-getting entertainment might reveal other benefits.

I have observed that similar beneficial effects occur for people also in role-playing scenarios, whether during group therapy where a person acts out an emotional situation, or during role-playing games such as the popular computer MUD’s and interactive communities where one is free to try out new personalities. A friend who is a Catholic priest once acknowledged how much he enjoyed getting to play an evil character in one of these role-playing games. Such entertainment provides a healthy and safe way to expand one’s emotional dynamic range.

Successful entertainment need not draw forth a roar of emotional response, a lot of wagging of reflectors, or a lot of pushing of buttons as in the interactive theaters coming soon from Sony. Nonetheless, even in the quietest symphony hall, a successful performer can sense how the audience is responding, and is, in turn, affected by their response.

Audience response could be captured by a variety of affective things – by cameras that looked at the audience, by active programs they hold in their hands, by chair arms and by floors that sense. Such affective sensors would add a new flavor of input to entertainment, providing dynamic forms that composers might weave into operas that interact with their audience.

For example, the floors in the intermission gathering spaces might be live compositions, waiting to sense the mood of the audience and amplify it with music. The environment itself might become a new musical instrument, perhaps like one of Machover’s hyperinstruments [64], but equipped to sense affect directly, augmenting the modes of expression available to the performer.

In general, audience appraisal would look not only at the people, but also at the media content, to distinguish, for example, the viewer’s sadness due to story content, e.g., the death of Bambi’s mom, and the viewer’s unhappiness due to other factors – possibly a degraded color channel, or garbled soundtrack. If the affective sensors were wearable, and perhaps seeing everything you see (See [65] for examples of these sensors), then they might correlate visual experiences with heart rate, respiration, and other forms of sentic modulation.

Affective computers provide a new set of tools, both for constructing new forms of entertainment, and for constructing new theories of what makes it succeed.

## 5.2 Expressive communication

*The power of essentic form in communicating and generating a sentic state is greater the more closely the form approaches the pure or ideal essentic form for that state.* – Seventh Principle of Sentic Communication [23]

Clynes [23] argues that music can be used to express emotion more finely than any language. But how can one master this finest form of expression? The master cellist Pablo Casals, advised his pupils repeatedly to “play naturally.” Clynes says he came to understand that this meant (1) to listen inwardly with utmost precision to the inner form of every musical sound, and then (2) to produce that form precisely. Clynes illustrates with the story of a young master cellist, at Casals’s house, playing the third movement of the Haydn cello concerto. All the attendees admired the grace with which he played – except Casals:

Casals listened intently. “No,” he said, and waved his hand with his familiar, definite gesture, “That must be graceful!” And then he played the same few bars – and it was graceful as though one had never heard grace before – a hundred times more graceful – so

that the cynicism melted in the hearts of the people who sat there and listened. [23]

Clynes attributes the power of Casals's performance to the purity and preciseness of the essential form. The purer the emotional state, the purer its expression and communication. In expression, teaches Clynes, faithfulness to the purest inner form produces the best results.

With affective recognition, the computer music teacher might not only try to hold your interest longer to help you learn more, but it might also give feedback as you develop preciseness of expression. This "emotional biofeedback," through measuring essential form, perhaps via finger pressure, foot pressure, or functions of inspiration and expiration as you breathe, could help you compare aspects of your performance that have never been measured or understood before.

Recently, Clynes [20] has made significant progress in this area, giving a user control over such expressive aspects as pulse, note shaping, vibrato, and timbre. Clynes recently conducted a "Musical Turing test"<sup>17</sup> to demonstrate the ability of his new "superconductor" tools. In this test, hundreds of people listened to seven performances of Mozart's sonata K330. Six of the performances were by famous pianists and one was by a computer. Most people could not discern which of the seven was the computer, and people who ranked the performances ranked the computer's as second or third on average. Clynes's computer's performances, which have played to the ears and hearts of many master musicians, demonstrate that we can identify and control meaningful expressive aspects of music, often called "the finest language of emotion."

### 5.2.1 Expressive mail and small talk

Although emotional states may be subtle in their modulation of expression, they are not subtle in their power to communicate, and correspondingly, to persuade. When sentic modulation is missing, misunderstandings occur. Consider the tremendous reliance of many people on email that is currently limited to text. Most people who use email have found themselves misunderstood at some point – their comments received with the wrong tone.

By necessity, email has had to develop its own set of symbols for encoding tone, "emoticons" such as :- ) and ;-( (turn your head to the left to recognize the smileys). However, these icons are limited; consequently, much affect-less email has resulted in a loss of productivity as people expend energy trying to undo misunderstandings, or as people expend time trying to word their email more carefully. When affect communication is most important, then person-to-person contact carries the most information; email presently carries the least.

Although it is often desirable to deliberately limit emotional expression, say, during card games or business negotiations, it is almost never desirable to be forced to do so by the available medium. To free email from this restriction, tools that recognize and express affect could augment text with other modes of expression such as voice, face, or potentially touch. In addition to intonation and facial expression recognition, current low-tech contact with keyboards could be augmented with simple attention to typing rhythm and pressure, as another key to affect. The new "ring mouse" could potentially pick up other features

<sup>17</sup>Although Turing eliminated sensory (auditory, visual, tactile, olfactory, taste) expressions from his test, one can imagine variations where each of these factors is included, e.g., music, faces, force feedback, electronic noses, and comestible compositions.

such as skin conductivity, temperature, and pulse, all observations which may be combined to identify emotional state. An "affective mouse" could collect finger pressure information to determine the values of your responses while you cruise the world wide web, picture libraries, or virtual museums. Although none of this information should be a forced broadcast; it could be an option available for the user, like the ability to use a videophone with or without the camera capped.

Encoding affective information along with a text message could tell an audio receiver what affect to use in reading the message to its recipient. Affective bits could also be used to set the expression on a "talking head" that delivers news to your low-bit-rate video receiver. Moreover, affective state information can be transmitted to the widest variety of media – visual, auditory, text – and decoded according to the abilities of that receiver. As technologists try to enable the fullest bandwidth human communication in limited bandwidth systems, it is essential to extract the key bits of information. In human communication, the affective bits are often key.

A helpful analogy is to consider the proficiency with which someone who knows you well can read your emotions. We are often wrong about first impressions, but as you get to know somebody you become increasingly adept at guessing what lies beneath their expression. A person you know well may walk in and say "good morning" with a particular tone that tells you they are in a hurry but happy to see you. Someone you don't know well may say "good morning" with the exact same tone, and you may interpret it quite differently.

What is really communicated when you say "good morning" or share other so-called "small talk" with someone you see regularly? The words have virtually no information content given their repeated occurrences. Might it be that the primary role of small talk is that of communicating affect?

Intelligent co-workers adapt their responses to you depending on the affect you express via small talk or other communication. Intelligent computers should be given the same opportunity, perhaps via an affectively spoken "good morning" exchanged during a login session.

### 5.2.2 Vocal/Facial Communication

Voice inflection is a subtle but powerful tool for communicating not only the morning's mood, but also important emotions such as anger, grace, and empathy. Inflection can also signal interest or disinterest. Without doubt, inflection adds flavor to our speech and content to its message. Even in telling a joke, everyone knows it's *how* you tell it that determines its success.

A variety of features of speech are modulated by emotion; these may be divided into the three categories of voice quality, utterance timing, and utterance pitch contour. (Murray and Arnott [26] provide a recent review of these features.) Although virtually no work seems to have been done on computer analysis of affect in voices, several features have been demonstrated for synthesizing intonation in computer-generated speech [66], [26].

With a suitable affective voice, computers can communicate in a more natural and social way with humans. Monotonous voice mail recordings and voice-reminder systems could vary their voices from day-to-day, like a human voice varies. Such variation would render them more pleasant, and could be used to more quickly flag important and urgent information.

Another form of affective computing that has already met with some success involves facial expression recognition. Faces appear to be the most important means for visual communication of emotion. Emotion-modeled faces can be used to give computers graphical faces which mimic the emotive expressions identified by Ekman [67], making the computer faces seem more

human. Several categories of human facial expression can be recognized by computers, both from still images [68] and from motion images [69], [67], the latter which is more reliable.

The encoding of facial expression parameters [67], [70] may also provide a simultaneously *efficient* and *meaningful* description for video compression, two attributes that satisfy important criteria for future coding systems [71]. Instead of sending over a new picture each time the person's face changes, you need only send their "basic emotion" faces once, and update with descriptions of their emotional state, and any slight variations.

### 5.3 Film/video

*A film is simply a series of emotions strung together with a plot...* though flippant, this thought is not far from the truth. It is the filmmaker's job to create moods in such a realistic manner that the audience will experience those same emotions enacted on the screen, and thus feel part of the experience. – Ian Maitland, Emmy Award winning director and editor

It is the job of the director to create onstage or onscreen, a mood that provokes a desired affect in the audience. A director or writer adjusts the characters in the movie or script until they "feel" right – until they communicate the intended message (and its emotions) to the user. An affective computer assistant might help the novice director discern what is not right when it doesn't feel right – for example, is the configuration of the set or the lighting in conflict with what typically evokes the intended mood?

Sometimes expressions of mood in film can be easily qualified – lighting from below to create an eerie effect, for example. However, determination of precisely what constitutes an essential form in different media is poorly understood in general. The forms by which emotions are communicated are an open area for research.

Despite a lack of understanding of how emotion is communicated, there is undoubtedly a power humans have to transfer genuine emotion – we sometimes say emotions are contagious. Clynes suggests that the purer the underlying essential form, the more powerful its communication. This power to purely transfer emotion exists not just from person to person, but also through external forms such as film.

#### 5.3.1 "Skip ahead to the interesting part"

My primary research for the last decade has focused on helping computers "see" as people see, with all the unknown and complicated aspects human perception entails. One of the newest applications of this research is the construction of tools that aid consumers and filmmakers in retrieving and editing video. Example goals are asking the computer to "find more shots like this" or to "fast forward to the dinosaur scene."

A much harder but related goal, is to teach a computer to "make a long story short." How does one summarize hours of video into a form pleasurable to browse? How do we teach the computer which parts look "best" to extract? Finding a set of rules that describe content for retrieving "more shots like this" is one difficulty, but finding the content that is "the most interesting" i.e., involving affect and attention, is a much greater challenge. These new challenges are ones which computer scientists are not equipped to address, but where cross-discipline efforts between cognitive science, emotion theory, and computer science are sorely needed.

The problem of locating a remembered scene, or an image with particularly interesting content, is also the problem of un-

derstanding causes of arousal, one of the key dimensions of affect. Arousal (excited/calm) has been found to be a better predictor of memory retention than valence (pleasure/displeasure) [72]. Image descriptions in Fig. 2 indicate associations of arousal with image content.

In fact, finding digital photographs having a particular "mood" was the most frequent request of advertising customers in a study of image retrieval made with the Kodak Picture Exchange [73]. Subject and action content, which were most frequently requested for editorial purposes, can also be a powerful contributor to mood in a photo [46].

We have recently built some of the first computer vision tools that enable computers to assist humans in annotating video, attaching descriptions to images that the person and computer both "see" [74]. Instead of the user tediously entering all the descriptions by hand, our algorithms learn which user-generated descriptions correspond to which image features, and then try to identify and label other "similar" content.<sup>18</sup>

Affective computing can be coupled with learning systems such as that of Minka and Picard [75], to begin to identify not only which content is most salient or interesting, but also which emotions tend to be evoked by the content. Successful learning algorithms for content-based similarity may also be able to learn examples of affect or mood similarity.

In fact, unlike trying to search for a shot that has a particular subject-action content, affective annotations, especially in terms of a few basic emotions or a few dimensions of emotion, could provide a relatively compact and salient index for retrieval of data. For example, people may tend to gasp at the same shots – "that guy is going to fall off the cliff!" Shots could be labeled (initially by a human) with descriptions such as "thrilling." The computer can later learn from the human which visual features best predict the most useful annotations.

For example, instead of annotating, "this is a sunny daytime shot of a student getting his diploma and jumping off the stage" the human might annotate "this shot of a student getting his diploma and jumping makes people grin." The latter is an affective annotation. Of course, although the latter example indicates a joyful shot for most viewers, it will not provoke a grin for everyone; an example is the mother whose son would have been at that graduation if he were not killed the week before. In other words, the cognitive and emotional state of the viewer interacts with what is perceived to produce the final affect.

These sorts of complicating factors will not be easy to address. Although affective annotations, like subject-action annotations, will not be universal, digitizing them will still help reduce time humans have to spend searching for the right scene. Both kinds of annotation are potentially powerful; we should be learning how both are perceived, and including them in digital audio and visual libraries.

### 5.4 Emotions in learning

Fascinating! – Spock, *Star Trek*

A learning episode might begin with curiosity and fascination. As the learning task increases in difficulty, one may experience confusion, frustration or anxiety. Learning may be abandoned because of these negative feelings. If the learner manages to avoid or proceed beyond these emotions then progress

<sup>18</sup>Computers have a hard time learning similarity, so this system tries to adapt to a user's ideas of similarity - whether perceptual, semantic, subjective, or otherwise.

may be rewarded with an “Aha!” and accompanying neuro-peptide rush. Even the “unemotional Spock” frequently exclaimed upon learning something new, that it was “fascinating!”

Dr. Barry Kort, a mentor of children exploring and constructing scientific worlds on the MUSE<sup>19</sup> and a volunteer for nearly a decade in the Discovery Room of the Boston Museum of Science, says that learning is the quintessential emotional experience [76]. Kort says his goal is to maximize intrigue – the fascination stage and to minimize anxiety.

Whatever her strategy, the good teacher detects important affective cues from the student and responds differently because of them. For example, the teacher might leave subtle hints or clues for the student to discover, thereby preserving the learner’s sense of self-propelled learning.

Whether the subject matter involves deliberate emotional expression such as music, or a “non-emotional” topic such as science, the teacher that attends to a student’s interest, pleasure, and distress, is perceived as more effective than the teacher that proceeds callously. The best human teachers know that frustration usually precedes quitting, and know how to skillfully redirect or motivate the pupil at such times. They get to know their student, and how much distress that student can withstand before learning breaks down.

Computers that cannot recognize human affect are severely handicapped. In contrast, with observations of your emotions, the computer could learn to respond to you more like the best human teachers, giving you one-on-one personalized guidance as you explore. Educational toys could have numerous learning strategies built in, changing their response as the child shows different levels of interest.

Affect has been largely ignored in theories of learning, perhaps because it is hard to measure. Like in all activities that demand mental performance, we know emotion is a determining factor. Hebb showed the classic inverted-U curve in [77] relating performance to arousal. Performance is lowest when the subject is awaking or when the subject is aroused to the point of emotional disturbance; performance is optimized at an intermediate state of arousal.

Now, suppose we could assess emotional communication during a learning episode – as a set of parameters, much like health researchers measure parameters such as heart-rate to develop better theories of athletic training. Perhaps there is some analog to be found in learning theory which corresponds to the rule-of-thumb in physical conditioning theory, “sustain 80% of maximum heart-rate for twenty minutes for optimal conditioning.” Although certainly a successful learning theory (for mind or for sports) would be more complicated than this, my point is that emotion variables could play a key role in developing new theories of learning. Such theories might begin to address issues such as not merely trying to optimize the student’s happiness, but orchestrating a balance between difficulty (frustration) and accomplishment (satisfaction).

Analyzing emotion both in the sender and receiver might also lead to progress in understanding the beneficial contagious effects of emotion – such as the life-changing impact of those special teachers who stir our interest in a topic.

When using a computer, people often find themselves trying to learn. Whether it is learning a new piece of software or directly using an educational program, the experiences are not generally considered pleasant. What if computer programs could, like the piano tutor above, pay attention to the user’s

affective expression. The online collection of this information could not only lead to more relevant feedback to the user, but could also be of great use in consumer satisfaction studies, not to mention in the development of more pleasurable and effective learning experiences.

Learning is almost always bi-directional; for example, the teacher usually learns things about the student’s attention and preferences during their interaction. Affective computer programs, particularly software agents which watch their users carefully, could begin to learn the human’s preferences, much like a trusted assistant. However, in the short term, like the dog walking the person and the new boots breaking in your feet, we will likely find ourselves doing as much adapting to the agents as they do to us. During this mutual learning process, it might be favorable if the agent at least paid attention to your frustration or satisfaction.

For example – the agent might notice our response to too much information as a function of valence (pleasure/displeasure) with the content. Too many news stories tailored to our interests might be annoying, and an occasional insertion of humor stories might lead to greater tolerance for the necessary but less pleasurable reading. Our tolerance may be described not only as a function of the day of week or time of day, but also as a function of baseline mood measured that morning. The agent, learning to distinguish which features of information best please the user while meeting his or her needs, could adjust itself appropriately. “User friendly” and “personal computing” might move closer to their stated meanings.

## 5.5 Affective environments

Sometimes people like a change of environment; sometimes it drives them crazy. These responses apply to all environments – not just your building, home, or office, but also your computer software environment with its “look and *feel*,” the interior of your automobile, and all the appliances with which you surround and augment yourself. What makes you prefer one environment to another?

Hooper [78] identified three kinds of responses to architecture, which I think hold true for all environments: (1) cognitive and perceptual – “hear/see,” (2) symbolic and inferential – “think/know,” and (3) affective and evaluative – “feel/like.” Perceptual computing (primarily computer vision and audition) and artificial intelligence have been largely concerned with measuring information in the first and second categories. Affective computing addresses the third.

In trying to understand what designs bring long-term satisfaction in his recent book *Buildings that Learn* [79], Stewart Brand emphasizes not the role of buildings as space, but their role in time. Brand applauds the architect who listens to and learns from post-occupancy surveys. But, he further cautions that because these are written or verbal reports, and the language of emotions is so inexact, these surveys are limited in their ability to capture what is really felt. Brand notes that surveys occur at a much later time than the actual experience, and hence may not recall what the visitors or inhabitants liked most.

In contrast, measuring sentic responses of people in the building could tell you how the customers feel when they walk in your bank vs. into the competitor’s bank, capturing those important “first impressions.” Surveys of willing newcomers who express their feelings when they enter your building for the first time might be recorded by an affective computer.

After being in a building awhile, your emotions in that space are no longer likely to be influenced by its structure, as that

<sup>19</sup>Email [oc@musenet.org](mailto:oc@musenet.org) or point your Gopher or Web browser at [cyberion.musenet.org](http://cyberion.musenet.org) for information how to connect.

has become predictable. Environmental factors such as temperature, lighting, sound, and decor – to the extent that they change – are more likely to affect you. “Alive rooms” or “alive furniture and appliances” that sense affective states could adjust factors such as lighting (natural or a variety of artificial choices) sound (background music selection, active noise cancellation) and temperature to either match or stimulate an appropriate mood. Your digital disc jockey might help suggest music of a particular mood. “Computer, please adjust the environment for a peaceful ambience at our party tonight.”

## 5.6 Aesthetic pleasure

Art does not think logically, or formulate a logic of behavior; it expresses its own postulate of faith. In science it is possible to substantiate the truth of one’s case and prove it logically to one’s opponents, in art it is impossible to convince anyone that you are right if the created images have left him cold, if they have failed to win him with a newly discovered truth about the world and about man, if in fact, face to face with the work, he was simply bored. – Andrey Tarkovsky [80]

As creation is related to the creator, so is the work of art related to the law inherent in it. The work grows in its own way, on the basis of common, universal rules, but it is not the rule, not universal *a priori*. The work is not law, it is above the law. – Paul Klee [81]

Psychology, sociology, ethnology, history, and other sciences have attempted to describe and explain artistic phenomena. Many have attempted to understand what constitutes beauty, and what leads to an aesthetic judgement. The elusiveness and complexity of aesthetics is due, in part, to the fact that affect plays a primary role.

Computers have been not only affect-blind, but aesthetically blind. Consider a scenario where a computer is assembling a presentation for you. In the not too distant future, the computer will be able to search digital libraries all over the world, looking for images and video clips with the content you request, e.g. “colorful scenes of Bill Gates with a cereal bowl.” Suppose it finds hundreds of shots that meet the requirements you gave it for content. What you would really like at that point is for it to narrow down the set to just the “good” ones to show you. How do you teach it what is “good?” Can something be measured in a picture, sculpture, building, piece of music, or flower arrangement that will indicate its beauty and appeal?

### 5.6.1 Hidden forms?

Clynes has suggested that essentic forms capture emotion in art. He has identified visual essentic forms in a number of great works of art – for example, the collapsed form of grief in the Pietà of Michelangelo (1499) and the curved essentic form of reverence in Giotto’s *The Epiphany* (1320). Clynes suggests that these visual forms, which match his measured finger-pressure forms, are indicative of a true internal essentic form. Moreover, shape is not the only parameter that could communicate this essentic form – color, texture, and other features may work collectively.

The viewpoint that we could find some combination of primitive elements in a picture that corresponds to an emotion is debated. On one hand, it seems that if you could find a basic essentic form in one image, that you could yank it out of that image, into some image-manipulation software, and construct a new image around it, one that does not communicate the same

emotion. What is the space of transformations that an essentic form could be put through, without it losing its identity and purity?

If such forms could be identified and measured, we could search the visual databases of the world-wide web for instances of them, and then test people’s responses to the images to see if the images communicate the same emotion. It seems it might be easy to find conflicting examples, especially across cultural, educational, and social strata. However, to my knowledge, no such investigation has been attempted yet.

Despite the lack of certainty about a visual form for emotion in pictures, there *is something* in a picture or piece of music that affects people emotionally; the question is what is the nature of this something? Note that even if such an investigation results in an ambiguous response, it would still not imply that Clynes’s hypothesized *internal* essentic forms do not exist, as seen in the following scenario:

One of Cytowic’s synesthetic patients saw lines projected in front of her when she heard music. Her favorite music makes the lines travel upward. If a computer could make the synesthetic associations she makes, then presumably it could help her find new music she would like, by looking for pieces where the lines went upward. However, although synesthetes have been found to make the same cross-modal associations throughout life, different synesthetes may have different associations. For an individual synesthete, rules might be discovered to predict these aesthetic feelings. Across synesthetes, the rules for the shape of the essentic form may be different. Nonetheless, there are consistent internal forms that an individual synesthete can use. What about internal forms for the rest of us?

Cytowic’s idea is that perception, in all of us, passes through the limbic system. Unlike most people, synesthetes can perceive the form while it is passing through this intermediate stage. Perhaps the limbic system is where Clynes’s hypothesized “essentic form” resides. Measurements of internal essentic forms may someday contribute to “objective” recognition of the aesthete. Just as lines going up and at an angle co-occurred with music the synesthetic woman liked, so certain essentic forms, and their purity, may be associated with preferences of other art forms.

With the rapid development of image and database query tools, we are entering a time where one could browse for examples of such forms; hence, this area is now more testable than ever before. But let’s again set aside the notion of trying to find a universal form, to consider a more personal scenario.

### 5.6.2 Personal taste

You are strolling past a store window and a garment catches your eye – “My friend would love that pattern!” you think. Later you look at a bunch of ties and mock – “How could anybody like these?”

People’s preferences differ wildly in clothing. They may reason about their taste along different levels – quality of the garment, its stitching and materials, its practicality or feel, its position in the fashion spectrum (style and price), and possibly even the reputation and expressive statement of its designer. A buyer knows, however, that these features cannot simply be optimized into a garment that all will find maximally appealing; it is absurd to assume the existence of a universal predictor of what is best.

Although you may or may not equate garments with art, an analogy exists between one’s aesthetic judgment in the two cases. Artwork is evaluated for its quality and materials, how well it will fit where you want to display it, its feel, its position in the world of art (style and price), its artist, and her or his

expressive intent. Just as for clothing, the finding of a universal aesthetic predictor may not be possible.

However, selecting something for someone you know well, something you think they would like, is commonly done. We not only recognize our own preferences, but we are often able to learn another's.

Moreover, clearly there is something in the appearance of the garment or artwork that influences our judgement. But what functions of appearance might enable the computer to reach the same judgement? Perhaps if the lines in that print were straighter, it would be too boring for you. On the other hand, you might treasure the bold lines on that table.

There are many problems with trying to find something to measure, be it in a sculpted style, painting, print, fabric, or room decor. Ordinary pixels and lines do not induce aesthetic feelings on their own, unless, perhaps it is a line of Klee, used to create an entire figure. Philosophers such as Langer have taken a hard stance in seeking to understand projective feeling in art:

There is, however, no basic vocabulary of lines and colors, or elementary tonal structures, or poetic phrases, with conventional emotive meanings, from which complex expressive forms, i.e., works of art, can be composed by rules of manipulation. [82]

Despite Langer's claim, neither do we experience aesthetic pleasure *without* the pixels, lines, notes and rhythms. Moreover, Clynes does seem to have found a set of mechanisms from which complex expressive forms can be produced, as evidenced in his musical Turing test; this is further collaborated with his recent studies indicating the significant role of composer's pulses in appreciation of music [83].

Just thinking of a magnificent painting or piece of music does not usually arouse the same emotions as when one is actually experiencing the work, but it may arouse similar, fainter emotions. Beethoven still composed some of the greatest music in the world after he could no longer hear. Aesthetic feelings appear to emerge from some combination of physical, perceptual, and cognitive forms.

Efforts to give computers recognition of what we think is beautiful should aid our understanding of this perplexing and universally important problem. Like all computer recognition problems, this one will probably require huge sets of examples of things that we do and do not like. Additionally, I expect it will be improved as the computer learns to directly incorporate affective feedback from us. The computer will need to infer which features are common to those examples that we like and distinguish these from features common to the negative examples.

With knowledge of affective preferences, the computer can cruise the networks at night, helping shop for clothes, furniture, wallpaper, music, gifts, artwork, and more. Online museum tours, that are becoming available on the world-wide web, could suggest to you additional collections, by guessing what you might like after observing your reaction to what you have already seen. Affective computers potentially provide more personal service, tailored to your ever-changing interests.

### 5.6.3 Design

You can not invent a design. You recognise it, in the fourth dimension. That is, with your blood and your bones, as well as with your eyes. – David Herbert Lawrence

Have you asked a designer how she arrived at the final design? Of course, there are design principles and constraints on

function that influenced her one way or the other. Such "laws" play an important role. However, none of them are inviolable. What seems to occur is a nearly ineffable recognition – a perceptual "aha!" when all is right.

Although we can measure qualities of objects, of the space between them, and of many components of design, we cannot predict how these alone will influence the experience of the observer. Design is not solely a rule-based process, and computer tools to assist with design only help explore a space of possibilities. Today's tools, e.g., in graphic design, incorporate principles of physics and computer vision to both judge and modify qualities such as balance, symmetry and disorder [84]. But the key missing objective of these systems is the goal of arousing an experience in the user – arousing to provoke attention, interest, memory, and new experiences. For this, the system must be able to recognize the user's affect dynamically, as the design is changed. (This assumes the user is a willing participant, not suppressing their feelings about the computer's design suggestions.)

Aesthetic success may be said to be communicated via emotion – you like something because it makes you feel good, or because you like to look at it, or it inspires you, or makes you think of something new; this brings you joy. A design solves a problem you have and you feel relief. Eventually, it brings you to a new state that is more satisfying than the one you were in.

Although the computer does not presently know how to lead a designer to this satisfied state, there is no reason it could not begin to store sentic responses, and gradually try to learn associations between these responses and the underlying design components. Sentic responses have the advantage of not having to be translated to language, which is an imperfect medium for reliably communicating feedback concerning design. In fact, frequently, it is precisely the sentic response that is targeted during design.

Affective computing will play a key role in gathering information for improving our aesthetic understanding, especially in areas such as entertainment and design.

## 5.7 Affective wearable computers

The most difficult thing is that affective states are not only the function of incoming sensory signals (i.e., visual, auditory etc.), but they are also the function of the knowledge/experiences of individuals, as well as of time. What you eat in the morning can influence the way you see a poster in the afternoon. What you read in tomorrow's newspaper may change the way you will feel about a magazine page you're just looking at now... – Suguru Ishizaki

The above peek into the unpredictable world of aesthetics emphasizes the need for computers which perceive what you perceive, and which recognize personal responses as you change them. In the most personal form, these are computers that could accompany you at all times.

The idea of wearing something that measures and communicates our mood is not new; the "mood rings" of the 70's are probably due for a fad re-run and mood shirts are supposedly available now locally. Of course these armpit heat-to-color transformers do not really measure mood. Nor do they compare to the clothing, jewelry, and accessories we could be wearing – lapel communicators, a watch that talks to a global network, a network interface that is woven comfortably into your jacket or vest, local memory and a microprocessor in your belt, a miniature videocamera and holographic display on your eyeglasses, and more.

Wearables may fulfill some of the dreams espoused by Clynnes when he coined the word “cyborg” [85]. Wearable computers can augment your memory (any computer accessible information available as you need it) [86] or your reality (zooms in when you need to see from the back of the room). Your wearable camera could recognize the face of the person walking up to you, and remind you of his or her name and where you last met. Signals can be passed from one wearable to the other through your conductive “BodyNet” [87]. A handshake could instantly pass to my online memory the information on your business card.<sup>20</sup> Note that these examples are not science fiction; all of these functions or their basic technologies have been realized in present research in the MIT Media Laboratory.

An Orson Scott Card science fiction novel [88] features a sentient being named Jane that speaks from a jewel in the ear of Ender, the hero of the story. To Jane, Ender is her brother, as well as dearest friend, lover, husband, father, and child. They keep no secrets from each other; she is fully aware of his mental world, and consequently, of his emotional world. Jane cruises the universe’s networks, scouting out information of importance for Ender. She reasons with him, plays with him, handles all his business, and ultimately persuades him to tackle a tremendous challenge. Jane is the ultimate affective and effective computer agent, living on the networks, and interacting with Ender through his wearable interface.

Although Jane is science fiction, agents that roam the networks and wireless wearables that communicate with the networks are current technology. Computers come standard with cameras and microphones, ready to see our facial expression and listen to our intonation. People who work with computers generally have more physical contact with computers than they have with people; computers are in a unique position to sense our affective state.

The bandwidth humans have for communicating thoughts and emotions to each other can also be available for communicating with computer agents. My wearable agent might be able to see your facial expression, hear your intonation, and recognize your speech and gestures. Your wearable might feel the changes in your skin conductivity and temperature, sense the pattern of your breathing, measure the change in your pulse, feel the lilt in your step, and more, all in an effort to better understand you. You could choose to whom your wearable would communicate these personal clues of your emotional state.

I want a mood ring that tells me my wife’s mood before I get home – Walter Bender

If we were willing to wear a pulse, respiration, or moisture monitor, the computer would have more access to our motor expression than most humans. This opens numerous new communication possibilities, such as the message (perhaps encrypted) to your spouse of how you are feeling as you head home from the office. The mood recognition might trigger an offer of information, such as the news (via the network) that the local florist just received a delivery of your spouse’s favorite protea. A mood detector might make suggestions about what foods to eat, so called “mood foods” [89], and collect information continuously through the diet, contributing to our ongoing understanding of biochemical influences on mood.

Affective wearables offer possibilities of new health and medical research opportunities and applications. Medical studies could move from measuring controlled situations in labs, to measuring more realistic situations in life. A jacket you choose to wear that senses your posture might gently remind you to

correct a bad habit after back surgery, perhaps by a subtle nudge in a helpful place. Wearables that measure other physiological responses can help you identify causes of stress and anxiety, and how well your body is responding to these.<sup>21</sup> Such devices might be connected to medical alert services, a community of friends and family, or perhaps just a private “slow-down and attend to what you’re doing” service, providing personal feedback for your private reflection – “I sense more joy in you tonight.”

With willing participants, and successful affective computing, the possibilities are limited only by our imagination. Affective wearables would be communication boosters, clarifying feelings, amplifying them when appropriate, and leading to imaginative new interactions and games. Wearables that detect your lack of interest during an important lecture might switch into a recording mode for you, taking notes while assuming that your mind is wandering. Games where players don affective computers might add points for courage. Your wearable might encourage you during a workout, “I sense anger reduction.” Wearables with networked agents might help people reach out to contact those who want to be contacted, not just based on common interests as expressed through internet news groups, but also through common mood. For example, it might recognize your emotional state could be improved by striking up a conversation with someone with common interests right now; and it might let you know who’s available that would enjoy this opportunity.

Of course, you could remap your affective processor to change your affective appearance, or to keep certain states private. In offices, one might wish to reveal only the states of no emotion, disgust, pleasure, and interest. Affective computing does not enforce emotion recognition or expression on anyone. Instead, it provides an opportunity for additional communication, one which can be used for both good and bad purposes, with hopeful emphasis on the former.

### 5.7.1 New data needed

Despite a number of significant efforts, emotion theory is in its infancy. People’s emotional patterns depend on the context in which they are elicited – and so far these have been limited to lab settings. Problems with studies of emotion in a lab setting (especially with interference from cognitive social rules) are well documented. The ideal study to aid the development of the theory of emotions is real-life observation, recently believed to be impossible [30].

However, as in the examples above, a wearable affective computer that attends to you during your waking hours could potentially notice what emotions you express, as well as a variety of conditioning factors such as what you eat, what you do, what you see, hear, etc. Computers excel at amassing information, and their ability to analyze and identify patterns is being improved rapidly. Given a willingness on the part of the wearer to share this information with researchers, a wealth of important data could be gathered for furthering theories of learning, intelligence, perception, diet, exercise, communication, mental health, and more.

## 6 Summary

Emotion was identified by Donald Norman in 1981 as one of the twelve major challenges for cognitive science [91]. In this paper I have argued that emotions can no longer be considered a luxury when studying essential rational cognitive processes;

<sup>20</sup>It could also pass along a virus.

<sup>21</sup>See [90] for a discussion of emotions and stress.

instead, recent neurological evidence indicates they are *necessary* not only in human creativity and intelligence, but also in rational human thinking and decision-making. I have suggested that if computers will ever interact naturally and intelligently with humans, then they need the ability to at least recognize and express affect.

Affective computing is a new area of research, with recent results primarily in the recognition and synthesis of facial expression, and the synthesis of voice inflection. However, these results are just the tip of the iceberg; a variety of physiological measurements are available which would yield clues to one's hidden affective state. Moreover, these states do not need to be universal in their expression for a computer to recognize them. I have proposed some possible models for the state identification, treating affect recognition as a dynamic pattern recognition problem. More research is needed to discover which of these tools, coupled with which measurements, both of the person and their environment, give reliable indicators of affect for an individual.

Given modest success recognizing affect, significant leaps in both theory and practice are possible. Affect plays a key role in understanding phenomena such as attention, memory, and aesthetics. I have described over fifty possible applications in learning, information retrieval, communications, entertainment, design, health, and human interaction where affective computing would be beneficial. In particular, with wearable computers that perceive context and environment as well as physiological information, there is the potential of gathering data for advances in cognitive and emotion theory, as well as for improving our understanding of factors that contribute to human health and well-being.

Although I have focused on computers that recognize and portray affect, I have also mentioned evidence for the importance of computers that would "have" emotion. Emotion is not only necessary for creative behavior in humans, but neurological studies indicate that decision-making without emotion can be just as impaired as decision-making with too much emotion. I have used this evidence to suggest that building computers that make intelligent decisions may require building computers that "have emotions."

I have also proposed a dilemma that arises if we choose to give computers emotions. Without emotion, computers are not likely to attain creative and intelligent behavior, but with too much emotion, we, their maker, may be eliminated by our creation. Although this scenario is far afield and included mostly as a worst-case possibility, it is important that researchers discuss potential hazards of affective computing together with its potential benefits.

I have suggested a wide range of benefits if we build computers that recognize and express affect. The challenge in building computers that not only recognize and express affect, but which *have emotion* and use it in learning and making decisions, is a challenge not merely of balance, but of wisdom and spirit. It is a direction into which we should proceed only with the utmost respect for humans, their thoughts, emotions, and freedom.

## Acknowledgments

It is not often that one reads books that redirect their course of research. I am indebted to Manfred Clynes, Richard E. Cytowic, and Antonio Damasio, whose pioneering studies into emotion argued persuasively for more attention on this topic, ultimately capturing my own. Thanks to Beth Callahan, Len Picard, and Ken Haase for enjoyable science fiction pointers and to Dave DeMaris and Manfred Clynes for reminding me of

the human qualities that, so importantly, should be asked of research like this. It is with gratitude I acknowledge Bil Burling, Manfred Clynes, Allan Collins, Richard E. Cytowic, Dave DeMaris, Peter Hart, Jennifer Healey, Mitch Kapor, Fanya Montalvo, Alex Pentland, Len Picard, Josh Wachman, and Alan Wexelblat for helpful and encouraging comments on early drafts of this manuscript.

## References

- [1] R. S. Lazarus, *Emotion & Adaptation*. New York, NY: Oxford University Press, 1991.
- [2] R. E. Cytowic, *The Man Who Tasted Shapes*. New York, NY: G. P. Putnam's Sons, 1993.
- [3] R. E. Cytowic, *Synesthesia : a union of the senses*. New York: Springer-Verlag, 1989.
- [4] C. E. Izard, "Four systems for emotion activation: Cognitive and noncognitive processes," *Psychological Review*, vol. 100, no. 1, pp. 68-90, 1993.
- [5] O. Kroeger and J. M. Thuesen, *Type Talk at Work*. New York: Delacorte Press, Bantam Doubleday Dell Publishing Group, Inc., 1992.
- [6] R. E. Cytowic, *The Neurological Side of Neuropsychology*. Cambridge, MA: MIT Press, 1995. To Appear.
- [7] A. R. Damasio, *Descartes' Error: Emotion, Reason, and the Human Brain*. New York, NY: Gosset/Putnam Press, 1994.
- [8] P. N. Johnson-Laird and E. Shafir, "The interaction between reasoning and decision making: an introduction," *Cognition*, vol. 49, pp. 1-9, 1993.
- [9] A. M. Turing, "Computing machinery and intelligence," *Mind*, vol. LIX, pp. 433-460, October 1950.
- [10] Aristotle, *The Rhetoric of Aristotle*. New York, NY: Appleton-Century-Crofts, 1960. An expanded translation with supplementary examples for students of composition and public speaking, by L. Cooper.
- [11] D. R. Hofstadter and D. C. Dennett, *The Mind's I*. Bantam Books, 1981.
- [12] D. Goleman, *Emotional Intelligence*. Bantam Books, 1995.
- [13] C. I. Nass, J. S. Steuer, and E. Tauber, "Computers are social actors," in *Proceeding of the CHI '94 Proceedings*, (Boston, MA), pp. 72-78, April 1994.
- [14] C. Nass and S. S. Sundar, "Is human-computer interaction social or parasocial?," *Submitted for Publication*, 1994. Stanford SRCT Paper #110.
- [15] S. R. Covey, *The Seven Habits of Highly Effective People*. New York: Fireside, Simon and Schuster, 1989.
- [16] N. Negroponte, *Being Digital*. New York: Alfred A. Knopf, 1995.
- [17] E. D. Scheirer, 1994. Personal Communication.
- [18] M. Lewis, "Self-conscious emotions," *American Scientist*, vol. 83, pp. 68-78, Jan.-Feb. 1995.
- [19] R. Plutchik and H. Kellerman, eds., *Emotion Theory, Research, and Experience*, vol. 1-5. Academic Press, 1980-1990. Series of selected papers.
- [20] M. Clynes, 1995. Personal Communication.



- [21] G. Duchenne, *The Mechanism of Human Facial Expression*. New York, NY: Cambridge University Press, 1990. Reprinting of original 1862 dissertation.
- [22] R. E. Cytowic, 1994. Personal Communication.
- [23] D. M. Clynes, *Sentics: The Touch of the Emotions*. Anchor Press/Doubleday, 1977.
- [24] R. S. Lazarus, A. D. Kanner, and S. Folkman, "Emotions: A cognitive-phenomenological analysis," in *Emotion Theory, Research, and Experience* (R. Plutchik and H. Kellerman, eds.), vol. 1, Theories of Emotion, Academic Press, 1980.
- [25] P. Ekman and W. Friesen, *Facial Action Coding System*. Consulting Psychologists Press, 1977.
- [26] I. R. Murray and J. L. Arnott, "Toward the simulation of emotion in synthetic speech: A review of the literature on human vocal emotion," *J. Acoust. Soc. Am.*, vol. 93, pp. 1097–1108, Feb. 1993.
- [27] K. Leidelmeijer, *Emotions: An experimental approach*. Tilburg University Press, 1991.
- [28] R. B. Zajonc, "On the primacy of affect," *American Psychologist*, vol. 39, pp. 117–123, Feb. 1984.
- [29] M. Lewis, "Ch 16: The emergence of human emotions," in *Handbook of Emotions* (M. Lewis and J. Haviland, eds.), (New York, NY), pp. 223–235, Guilford Press, 1993.
- [30] H. G. Wallbott and K. R. Scherer, "Assessing emotion by questionnaire," in *Emotion Theory, Research, and Experience* (R. Plutchik and H. Kellerman, eds.), vol. 4, The Measurement of Emotions, Academic Press, 1989.
- [31] A. Wierzbicka, "Defining emotion concepts," *Cognitive Science*, vol. 16, pp. 539–581, Oct.-Dec. 1992.
- [32] J. D. Laird, J. J. Wagener, M. Halal, and M. Szegda, "Remembering what you feel: Effects of emotion on memory," *Journal of Personality and Social Psychology*, vol. 42, no. 4, pp. 646–657, 1982.
- [33] D. Gelernter, *The Muse in the Machine*. Ontario: The Free Press, Macmillan, Inc., 1994.
- [34] M. Minsky, *The Society of Mind*. New York, NY: Simon & Schuster, 1985.
- [35] W. R. Hess, *The functional organization of the diencephalon*. New York: Grune & Stratton, 1957.
- [36] K. Levitt, 1995. Personal Communication.
- [37] M. Chekhov, *On the Technique of Acting*. New York: HarperCollins Pub., 1991. Chekhov, (1891-1955), Preface and Afterword by M. Powers.
- [38] R. Plutchik, "A general psychoevolutionary theory of emotion," in *Emotion Theory, Research, and Experience* (R. Plutchik and H. Kellerman, eds.), vol. 1, Theories of Emotion, Academic Press, 1980.
- [39] A. Ortony, G. L. Clore, and A. Collins, *The Cognitive Structure of Emotions*. Cambridge, MA: Cambridge University Press, 1988.
- [40] N. L. Stein and K. Oatley, eds., *Basic Emotions*. Hove, UK: Lawrence Erlbaum Associates, 1992. Book is a special double issue of the journal *Cognition and Emotion*, Vol. 6, No. 3 & 4, 1992.
- [41] L. R. Rabiner and B. H. Juang, "An introduction to hidden Markov models," *IEEE ASSP Magazine*, pp. 4–16, Jan. 1986.
- [42] K. Popat and R. W. Picard, "Novel cluster-based probability models for texture synthesis, classification, and compression," in *Proc. SPIE Visual Communication and Image Proc.*, vol. 2094, (Boston), pp. 756–768, Nov. 1993.
- [43] A. Sherstinsky and R. W. Picard, "Orientation-sensitive image processing with M-Lattice: A novel non-linear dynamical system," in *IEEE First Int. Conf. on Image Proc.*, vol. III, pp. 152–156, Nov. 1994.
- [44] L. A. Camras, "Expressive development and basic emotions," *Cognition and Emotion*, vol. 6, no. 3 and 4, 1992.
- [45] W. J. Freeman, *Societies of Brains. A Study in the Neuroscience of Love and Hate*. Hillsdale, NJ: Lawrence Erlbaum Associates, Publishers, 1995.
- [46] P. J. Lang, "The emotion probe: Studies of motivation and attention," *American Psychologist*, vol. 50, no. 5, pp. 372–385, 1995.
- [47] R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*. Wiley-Interscience, 1973.
- [48] C. W. Therrien, *Decision Estimation and Classification*. New York: John Wiley and Sons, Inc., 1989.
- [49] P. Werbos, "The brain as a neurocontroller: New hypotheses and new experimental possibilities," in *Origins: Brain and Self-Organization* (K. H. Pribram, ed.), Erlbaum, 1994.
- [50] E. J. Sondik, "The optimal control of partially observable Markov processes over the infinite horizon: Discounted costs," *Operations Research*, vol. 26, pp. 282–304, March-April 1978.
- [51] W. S. Lovejoy, "A survey of algorithmic methods for partially observed Markov decision processes," *Annals of Operations Research*, vol. 28, pp. 47–66, 1991.
- [52] C. C. White, III, "A survey of solution techniques for the partially observed Markov decision process," *Annals of Operations Research*, vol. 32, pp. 215–230, 1991.
- [53] T. Darrell, "Interactive vision using hidden state decision processes." PhD Thesis Proposal, 1995.
- [54] R. Pfeifer, "Artificial intelligence models of emotion," in *Cognitive Perspectives on Emotion and Motivation* (V. Hamilton, G. H. Bower, and N. H. Frijda, eds.), vol. 44 of *Series D: Behavioural and Social Sciences*, (Netherlands), pp. 287–320, Kluwer Academic Publishers, 1988.
- [55] P. Maes, T. Darrell, B. Blumberg, and A. Pentland, "The ALIVE system: Full-body interaction with autonomous agents," *Proceedings of Computer Animation*, Apr. 1995.
- [56] N. Alm, I. R. Murray, J. L. Arnott, and A. F. Newell, "Pragmatics and affect in a communication system for non-speakers," *Journal of the American Voice*, vol. 13, pp. 1–15, March 1993. I/O Society Special Issue: People with Disabilities.
- [57] D. Dennett, "Where am I?," in *Brainstorms: Philosophical Essays on Mind and Psychology*, (Montgomery, VT), Bradford Books, 1978.
- [58] I. Asimov, *The Bicentennial Man and Other Stories*. Garden City, NY: Doubleday Science Fiction, 1976.
- [59] A. C. Clarke, *2001 A Space Odyssey*. New York, NY: New American Library, Inc., 1968. based on the 1965 screenplay by S. Kubrick and A. C. Clarke.

- [60] D. R. Hofstadter, *Godel, Escher, Bach: an Eternal Golden Braid*. New York, NY: Basic Books, Inc. Publishers, 1979.
- [61] G. W. V. Leibniz, *Monadology and Other Philosophical Essays*. Indianapolis: The Bobbs-Merrill Company, Inc., 1965. Essay: Critical Remarks Concerning the General Part of Descartes' Principles (1692), Translated by: P. Schrecker and A. M. Schrecker.
- [62] A. H. (1894-1963), *Brave new world & Brave new world revisited*. New York, NY: Harper & Row, 1965.
- [63] S. Tucker, 1995. Media Lab Spring Colloquium.
- [64] T. Machover, "Hyperinstruments: A composer's approach to the evolution of intelligent musical instruments," *CyberArts*, pp. 67-76, Jan. 1992.
- [65] S. Mann, "'See the world through my eyes,' a wearable wireless camera," 1995. <http://www-white.media.mit.edu/~steve/netcam.html>.
- [66] J. E. Cahn, "The generation of affect in synthesized speech," *Journal of the American Voice I/O Society*, vol. 8, July 1990.
- [67] I. A. Essa, *Analysis, Interpretation and Synthesis of Facial Expressions*. PhD thesis, MIT Media Lab, Cambridge, MA, Feb. 1995.
- [68] G. W. Cottrell and J. Metcalfe, "EMPATH: Face, emotion, and gender recognition using holons," in *Neural Information Processing Systems* (R. P. Lippmann, J. E. Moody, D. S. Touretzky, and B. M. Spatz, eds.), vol. 3, (San Mateo, CA), pp. 564-571, Morgan Kaufmann Publishers, Inc., 1991.
- [69] Y. Yacoob and L. Davis, "Computing spatio-temporal representations of human faces," in *Computer Vision and Pattern Recognition Conference*, pp. 70-75, IEEE Computer Society, 1994.
- [70] S. Morishima, "Emotion model - A criterion for recognition, synthesis and compression of face and emotion." 1995 Int. Workshop on Face and Gesture Recognition, to Appear.
- [71] R. W. Picard, "Content access for image/video coding: 'The Fourth Criterion'," Tech. Rep. 295, MIT Media Lab, Perceptual Computing, Cambridge, MA, 1994. MPEG Doc. 127, Lausanne, 1995.
- [72] B. Reeves, "Personal communication, MIT Media Lab Colloquium," 1995.
- [73] D. Romer, "The Kodak picture exchange," April 1995. seminar at MIT Media Lab.
- [74] R. W. Picard and T. P. Minka, "Vision texture for annotation," *Journal of Multimedia Systems*, vol. 3, pp. 3-14, 1995.
- [75] T. P. Minka and R. W. Picard, "Interactive learning using a 'society of models'," *Submitted for Publication*, 1995. Also appears as MIT Media Lab Perceptual Computing TR#349.
- [76] B. W. Kort, 1995. Personal Communication.
- [77] D. O. Hebb, *A Textbook of Psychology*. Philadelphia: W. B. Saunders Co., 1966.
- [78] K. Hooper, "Perceptual aspects of architecture," in *Handbook of Perception: Perceptual Ecology* (E. C. Carterette and M. P. Friedman, eds.), vol. X, (New York, NY), Academic Press, 1978.
- [79] S. Brand, *How buildings learn: what happens after they're built*. New York, NY: Viking Press, 1994.
- [80] A. Tarkovsky, *Sculpting in Time: Reflections on the Cinema*. London: Faber and Faber, 1989. ed. by K. Hunter-Blair.
- [81] P. Klee, *The Thinking Eye*. New York, NY: George Wittenborn, 1961. Edited by Jurg Spiller, Translated by Ralph Manheim from the German 'Das bildnerische Denken'.
- [82] S. K. Langer, *Mind: An Essay on Human Feeling*, vol. 1. Baltimore: The Johns Hopkins Press, 1967.
- [83] M. Clynes, "Microstructural musical linguistics: composers' pulses are liked most by the best musicians," *Cognition*, vol. 55, pp. 269-310, 1995.
- [84] X. Yang, "Visual balance: The tightrope of computer generated layout," Master's thesis, MIT, Cambridge, MA, June 1995.
- [85] M. Clynes and N. S. Kline, "Cyborgs and space," *Astrophysics*, vol. 14, pp. 26-27 and 74-76, Sept. 1960.
- [86] T. E. Starner, "Wearable computing," Perceptual Computing Group, Media Lab 318, MIT, Cambridge, MA, 1995.
- [87] N. Gershenfeld and M. Hawley, 1994. Personal Communication.
- [88] O. S. Card, *Speaker for the Dead*. New York, NY: Tom Doherty Associates, Inc., 1986.
- [89] J. J. Wurtman, *Managing your mind and mood through food*. New York: Rawson Associates, 1986.
- [90] G. Mandler, *Mind and Body: Psychology of Emotion and Stress*. New York, NY: W. W. Norton & Company, 1984.
- [91] D. A. Norman, "Twelve issues for cognitive science," in *Perspectives on Cognitive Science* (D. A. Norman, ed.), (Hillsdale, NJ), pp. 265-295, Erlbaum, 1981.