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

Computers are beginning to acquire the ability to express and recognize affect, and may soon be given the ability to "have emotions." The essential role of emotion in both human cognition and perception, as demonstrated by recent neurological studies, indicates that affective computers should not only provide better performance in assisting humans, but also might enhance computers' abilities to make decisions. This paper presents and discusses key issues in "affective computing," computing that relates to, arises from, or influences emotions. Models are suggested for computer recognition of human emotion, and new applications are presented for computer-assisted learning, perceptual information retrieval, arts and entertainment, and human health and interaction. Affective computing, coupled with new wearable computers, will also provide the ability to gather new data necessary for advances in emotion and cognition theory.

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 has been prompted by fear. Nonetheless, the role of emotions is marginalized at best.

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

In this essay I will submit for discussion a set of ideas on what I call "affective computing," computing that relates to, arises from, or influences emotions. This will need some further clarification which I shall attempt below. I should say up front that I am *not* proposing the pursuit of computerized cingulotomies[1] or even into the business of building "emotional computers".

Nor will I propose answers to the difficult and intriguing questions, "what are emotions?" "what causes them?" and "why do we have them?"[2]

Instead, by a variety of short scenarios, I will define important issues in affective computing. I will suggest models for affect recognition, and present my ideas for new applications of affective computing to computer-assisted learning, perceptual information retrieval, arts and entertainment, and human health and interaction. I also describe how advances in affective computing, especially combined with wearable computers, can help advance emotion and cognition theory. First, let us begin with a brief scenario.

1.1 Songs vs. laws

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

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 this 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 think "I'll kill him when he gets here." When he finally shows, after you have nearly given up your last promise, how do you respond? Whether you are ready to greet him with rage or relief, does not his expression throw your switch? Your response swings if he arrives appearing inconsiderately carefree, or with woeful countenance. This response greatly affects what happens next.

Emotion pulls the levers of our lives, whether it be by the song in our heart, or the curiosity that drives our scientific inquiry. Rehabilitation counselors, pastors, parents, and to some extent, politicians, know that it is not laws that exert the greatest influence on people, but 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, the cultural "tune." I'm *not* suggesting we do away with laws, or even the rules (albeit brittle) that constitute rule-based artificial intelligence systems; Rather, I am saying that the laws and rules are not the most important part in human behavior. Nor do they appear to play the primary role in perception, as illustrated in the next scenario.

---

[1]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.

[2]For a list of some open questions in the theory of emotion, see Lazarus [1].

1.2 Limbic perception

> "Oh, dear," he said, slurping a spoonful, "there aren't enough points on the chicken." – Michael Watson, in [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 not walls between what is seen, felt, touched, smelled, and tasted. However, the neurological explanation for this perceptual phenomenon is not merely "crossed-wires."

The neurologist Cytowic has studied the neurophysiological aspects of synesthetic experience [2]. The cortex, usually regarded as the home of sensory perception, is expected to show increased activity during synesthetic experiences, where patients experience external and involuntary sensations somewhat like a cross-wiring of the senses – for example certain smells may elicit seeing strong colors. One would expect that during this heightened sensory experience, there would be an increase in cortical activity. However, during synesthesia, there is actually a collapse of cortical metabolism.[3]

Cytowic's studies point to a corresponding increase in activity in the limbic system, which lies physically between the brain stem and the two hemispheres of the cortex, and which has traditionally been assumed to play a less influential role than the cortex, which lies "above" it. The limbic system is the seat of memory, attention, and emotion. The studies during episodes of synesthesia indicate that the limbic system plays a central role in sensory perception.

Izard, in an excellent treatise on emotion theory [3], also describes emotion as a motivating and guiding force in perception and attention. Leidelmeijer [4] goes a step further in relating emotions and perception:

> Once the emotion process is initiated, deliberate cognitive processing and physiological activity may influence the emotional experience, but the generation of emotion itself is hypothesized to be a perceptual process.

In fact, there is a reciprocal relationship between the cortex and limbic system; they function in a closely intertwined manner. However, the discovery of the limbic role in perception, and of substantially more connections from the limbic system to the cortex, suggests that the limbic influence may be the greater. Cytowic is not the first to argue that the neurophysiological influence of emotion is greater than that of objective reason. The topic fills philosophy books and fuels hot debates. Often the limbic role is subtle enough to be consciously ignored – we say "Sorry, I guess I wasn't thinking" but not "Sorry, I wasn't feeling." Notwithstanding, the limbic system is a crucial player in our mental activity. If the limbic system is not directing the show, then it is at least a rebellious actor that has won the hearts of the audience.

1.3 Thinking–feeling axis

The limbic role is sometimes considered to be antithetical to thinking. The popular Myers-Briggs Type Indicator, has "thinking vs. feeling" as two endpoints of one of its axes for qualifying personality. People are quick to polarize thoughts and feelings as if they were opposites. But, neurologically, the brain draws no hard line between thinking and feeling:

> ...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, multisensory input is salient or not. [5].

1.3.1 Nonlimbic emotion and decision making

Although the limbic brain is the "home base" of emotion, it is not the only part of the brain engaged in the experience of emotion. The neurologist Damasio, in his book, *Descartes' Error* [6] identifies several non-limbic regions which affect emotion, and surprisingly, its role in reason.

Most adults know that too much emotion can wreak havoc on reasoning, but less known is the recent evidence that *too little* emotion can also wreak havoc on reasoning. Years of studies on patients with frontal-lobe disorders indicate that impaired ability to feel yields impaired ability to make decisions; in other words, there is no "pure reason" [6]. Emotions are vital for us to function as rational decision-making human beings.

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 [7]. 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 [6]. Damasio's theory is that emotion plays a biasing role in decision-making. One might say emotion wards off an infinite logical search. How do you decide how to proceed given scientific evidence? There is not time to consider *every possible* logical path.

I must emphasize 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. However, we must not marginalize the role of the "songs." The neurological evidence indicates emotions are not a luxury; they are essential for rational human performance. Belief in "pure reason" is a logical howler.

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 should it both think and feel? Let us consider the classic test of a thinking machine.

1.3.2 The Turing test

The Turing test examines if, in a conversation between a human and a computer, the human cannot tell if the computer's replies are being generated by a human (say, behind a curtain) or by the machine. The Turing test is considered a test of whether or not a machine can "think," in the truest sense of duplicating mental activity – both cortical and limbic. One might converse with the computer 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. Although the Turing test is designed to take place communicating only via text, so that sensory expression (e.g., voice intonation and facial expression) does not play a role, emotions can still be perceived in text, and can still be elicited by its content and form [8]. Clearly, a machine will not pass the Turing test unless it is also capable of perceiving and expressing emotions.

Nass *et al.* have recently conducted a number of classical tests of human social interaction, substituting computers into

---

[3]Measured by the Oberist-Ketty xenon technique.

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. In these experiments they have repeatedly found that the results of the human-human studies still hold. Their conclusion is that individuals' interactions with computers are inherently natural and social [9]. Since emotion communication is natural between people, we should interact more naturally with computers that recognize and express affect.

Negroponte reminds us that even a puppy can tell when you are angry with it [10]. Computers should have at least this much affect recognition. Let's consider a scenario, you are going for some private piano lessons from a computer.

1.4 The effective piano teacher

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. [11]. This goal contains many challenges, one of the hardest which involves expression recognition, distilling the essential pitches of the music from its expression. Recognizing and interpreting affect in musical expression is very important, and I'll return to it again below. But first, there is an even more important component. This component is present in all teaching and learning interactions.

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 feelings. Imagine it has the ability to distinguish even the three emotions we were all born with – interest, pleasure, and distress [12].[4]

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 might slow things down and give you encouraging suggestions. Detecting user distress, without the user making mechanical playing errors, might signify a moving requiem, a stuck piano key, or the need to prompt for more information.

Whether the subject matter involves deliberate emotional expression such as music, or a "non-emotional" topic such as science, the teaching system still tries to maximize pleasure and interest, while minimizing distress. The best human teachers know that frustration usually precedes quitting, and know how to skillfully redirect the pupil at such times. With observations of your emotions, the computer teacher could respond to you more like the best human teachers, giving you one-on-one personalized guidance as you explore.

1.4.1 Quintessential emotional experience

Fascinating! – Spock, *Star Trek*

Dr. Barry Kort, a mentor of children exploring and constructing scientific worlds on the MUSE[5] 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 [13]. 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 may be rewarded with an "Aha!" and accompanying neuropeptide rush. Kort says his goal is to maximize intrigue – the "fascinating" stage and to minimize anxiety. The good teacher detects these important cues, and responds appropriately. 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.

Enthusiasm is contagious in learning. The teacher who expresses excitement about the subject matter can often stir up similar feelings in the student. Thus, in the above piano-teaching scenario, in addition to the emotional expression in the music, and the emotional state of the student, there is also the affect expressed by the teacher.

Computer teaching and learning systems abound, with interface agents perhaps providing the most active research area for computer learning. Interface agents are expected to be able to learn our 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, learning will be two-way. We will find ourselves doing as much adapting to the agents as they do to us. During this mutual learning process, wouldn't it be preferable if the agent paid attention to whether we were getting frustrated with it?

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; but some days, there can't be too many humor stories. Our tolerance may be described as a function not only of the day of week or time of day, but also of our mood. 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" would move closer to their true meanings.

The above scenario raises the issue of observing not just someone's emotional expression, but also their emotional state. Is it some metaphysical sixth sense with which we discern unvocalized feelings of others? If so, then we can not address this scientifically, and I am not interested in pursuing it. But clearly there are ways we discern emotion – through voice, facial expression, and other aspects of our so-called body language. Moreover, there is evidence that we can build systems that begin to identify both emotional expression, and its generating state.

2 Sentic[6] 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 [14]

---

[4]This view of [12] is not unchallenged; facial expression in the womb, as well as on newborns, has yet to receive an explanation with which all scientists agree.

[5]Point your Gopher or Web browser at cyberion.musenet.org, or email oc@musenet.org for information how to connect.

[6]"Sentic" is from the Latin *sentire*, the root of the words "sentiment" and "sensation."

2.1 Poker face, poker body?

The level of control involved in perfecting one's "poker face" is praised by society. 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 the flood. I spot the lilt in your walk today and therefore expect you are in a good mood. Although I might successfully conceal the nervousness in my voice, I am not able to suppress it throughout my body; you might find evidence if you grab my clammy hand.

Although debate persists about the nature of the coupling between emotion and physiological response, most writers accept a physiological component in their definitions of emotion. Lazarus *et al.* [15] 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.

Clynes exploits the physiological component of emotion supremely in the provocative book, Sentics. He formulates seven principles for sentic (emotional) communication, which pertain to "sentic states," a description given by Clynes to emotional states, largely to avoid the negative connotations associated with "emotional." Clynes emphasizes that emotions modulate our physical communication; the motor system acts as a carrier for communicating our sentic state.

2.2 Visceral and cognitive emotions

The debate: "Precisely what are the cognitive, physical, and other aspects of emotion?" remains unanswered by laboratory studies.[7] Attempts to understand the components of emotion and its generation are complicated by many factors, one of which concerns the problem of describing emotions. Wallbott and Scherer [17] emphasize the problems in attaching adjectives to emotions, as well as the well-known problems of interference due to social pressures and expectations, such as the social "display rules" found by the psychologist Ekman, in his studies of facial expression. For example, some people might not feel appropriate expressing disgust during a laboratory study.

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 have argued that cognitive appraisal is a necessary precondition for affective arousal. However, this view is refuted by the large amount of empirical evidence that affect can also be aroused without cognitive appraisal [18], [3].

A helpful distinction for sorting the non-cognitively-generated and cognitively-generated emotions is made by Damasio [6] who distinguishes between "primary" and "secondary" emotions. Note that Damasio's use of "primary" is more specific than the usage of "primary emotions" in the emotion literature. Damasio's idea is that there are certain features of stimuli in the world that we respond to emotionally first, and which activate a corresponding set of feelings (and cognitive state) secondarily. Such emotions are "primary" and reside in the limbic system. He 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. For secondary emotions, the

---

[7]It is beyond the scope of this essay to overview the extensive literature; I will refer the reader instead to the carefully assembled collections of Plutchik and Kellerman [16]. The quotes I use, largely from these collections, are referenced to encourage the reader to revisit them in their original context.

limbic structures are not sufficient; prefrontal and somatosensory cortices are also involved.

The body usually responds to emotion, although James's 1890 view of this response *being* the emotion is not accepted today [3]. Studies 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 (autonomic nervous system) 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" the emotion. Similar complications arise when trying to identify physiological responses which co-occur with emotional states (e.g., heart rate also increases when exercising). Leidelmeijer overviews several conflicting studies in [4], reminding us that a specific situation is not equally emotional for all people and an individual will not be equally emotional in all situations.

2.2.1 No one can read your mind

The issue for affective computing comes down to the following: We 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). However, we can measure physiological responses (facial expression, and more, below) which often arise during expression of emotion. We should at least be able to measure physiologically those emotions which are already manifest to others. How consistent will these measurements be when it comes to identifying the corresponding sentic state?

Leidelmeijer [4] discusses the evidence both for and against universal autonomic patterning. One of the outstanding problems is that sometimes different individuals exhibit different physiological responses to the same emotional state. However, this argument fails in the same way as the argument for "speaker-independent" speech recognition systems, where one tries to decouple the semantics of what is said from its physical expression. Although it would be a terrific accomplishment to solve this universal recognition problem, it is unnecessary, as Negroponte pointed out years ago. If the problem can be solved in a speaker-dependent way, so that your computer can understand you, then your computer can translate to the rest of the world.

The experiments in identifying underlying sentic state from observations of physical expression only need to demonstrate consistent patterning for an individual in a given perceivable context. The individual's personal computer can acquire ambient perceptual and contextual information (e.g., see if you're climbing stairs, detect if the room temperature changed, etc.) to identify autonomic emotional responses conditioned on perceivable non-emotional factors. Perceivable context should include not only physical milieu, but also cognitive milieu – for example, the information that this person has a lot invested in the stock market, and may therefore feel unusually anxious as its index drops. The priorities of your agent could shift with your affective state.

2.2.2 Emotional experience, expression, and state

The experience arises out of a Gestalt-like concatenation of two major components: visceral arousal and cognitive evaluation... What we observe are symptoms of that inferred emotional state – symptoms that range from language to action, from visceral symptoms to facial ones, from tender words to vio-

lent actions. From these symptoms, together with an understanding of the prior state of the world and the individual's cognitions, we infer a private emotional state. – George Mandler [19]

Let me briefly clarify some terminology – especially to distinguish emotional experience, expression, and state. I use *sentic state, emotional state,* and *affective state* interchangeably. These refer to your dynamic state when you experience an emotion. All you consciously perceive in such a state is referred to as your emotional *experience.* Some authors equate this experience with "emotional feelings." Your emotional state cannot be directly observed by another person. What you reveal, either voluntarily or not, is your emotional *expression,* or "symptoms" in Mandler's quote. This expression through the motor system, or "sentic modulation" helps others guess your emotional state.

When subjects are told to experience a particular emotional state, or when such a state is encouraged or induced (perhaps by listening to a story or watching a film), then they may or may not *express* their emotional state. If asked explicitly to express it, then autonomic responses are usually enhanced. Deliberate expression makes it easier to infer the underlying emotional state.

2.2.3 "Get that look off your face"

Facial expressions are one of the two most widely acknowledged forms of sentic modulation. Duchenne de Boulonge, in his 1862 thesis (republished in [20]) identified completely independent expressive face muscles, such as the muscle of attention, muscle of lust, muscle of disdain or doubt, and muscle of joy. Most present day attempts to recognize facial expression are based on the subsequent Facial Action Coding System of psychologist Paul Ekman [21], which provides mappings between measurable muscles and an emotion space.

Emotion-modeled faces can be used to give computers graphical faces which mimic these precise expressions identified by Ekman [22], making the computer faces seem more human. Yacoob and Davis [23] and Essa and Pentland [22] have also shown that several categories of human facial expression can be recognized by computers. The encoding of facial expression parameters [22], [24] may also provide a simultaneously efficient and *meaningful* description for image compression, two attributes that satisfy important criteria for future image coding systems [25]. Instead of sending over a new picture every 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.

2.2.4 It's not what she said, but how she said it

The second widely acknowledged form of sentic modulation is in voice. 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, who we assume can't 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 is greater than the words spoken.

Cahn [27] has demonstrated various features that can be adjusted during voice synthesis to control the affect expressed in a computer's voice. Control over affect in synthetic speech is also an important ability for speaking-impaired people who rely upon voice synthesizers to communicate verbally [28].

A variety of features of speech are modulated by emotion; Murray and Arnott [26] provide a recent review of these features, which they divide into the three categories of voice quality, utterance timing, and utterance pitch contour. They discuss how these parameters might be manipulated to give computers the ability to speak with affect.

2.2.5 Beyond face and voice

Other forms of sentic modulation have been explored by Clynes in [14]. 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: a smile, tone of voice, etc.

The motor output explored most carefully by Clynes is the transient pressure of a finger during voluntary sentic expression. This finger pressure response has been measured for thousands of people, and found to be not only repeatable, but to reveal distinct traces of "essentic form" for states such as no emotion, anger, hate, grief, love, joy, sex, and reverence [14]. 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 essentic forms.

There are many physiological responses which vary with time and which might potentially be combined to assist recognition of sentic states. These include heart rate, diastolic and systolic blood pressure, pulse, pupillary dilation, respiration, skin conductance and temperature. Some of these are revisited below with "affective wearable computers."

2.2.6 Sentic hotline

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, even 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 [29].

Although we are not all experts at reading faces, and comedians and actors can be quite good at feigning emotions, it is claimed that the attentive observer is always able to recognize a false smile [20]. [8] This is consistent with the findings of Duchenne de Boulonge 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 unmasks a false friend. [20]

The neurology literature also indicates 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 [30]. 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 [20].

---

[8]This view is debated, e.g., by [1] who claims that all phenomena that change with emotion also change for other reasons, but these claims are unproven.

2.2.7 Inducement of sentic states

> 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 [31]

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 [8] to instructing speakers how to arouse different emotions in his or her audience. Although inducement of emotions may be deliberate, it seems we, the receiver, often enjoy its effects. Certainly, we enjoy picking a stimulus such as music that will affect our mood in a particular way. We tend to believe that we are also free to choose our response to the stimulus. An open question is, are we always free to do so? In other words, can some part of our nervous system be externally 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. Izard overviews some of the controversial evidence for and against these claims [3]. One of the interesting questions relates to involuntary movements, such as through eye saccades [32] – can you view imagery that will cause your eyes to move in such a way that their movement induces a corresponding sentic state? Although the answer is unknown, the answers to questions like this may hinge on only a slight willingness[9] to be open to inducement. This question may evoke disturbing thoughts of potentially harmful mind and mood control; or potentially beneficial, depending on how it is understood, and with what goals it is used.

2.3 Sentic state pattern recognition

What thoughts and feelings are expressed, are communicated through words, gesture, music, and other forms of expression – all imperfect, bandlimited modes. 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 feelings.

However, the scientific recognition of affective state does appear 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 during voluntary expression, but may also be found to be useful during involuntary 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 an underlying emotion or combination of 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 imaging and scanning techniques, might well be the largest "inverse problem" imaginable. In contrast, for emotion recognition, a relatively small number of simplifying categories for emotions have been commonly proposed.

2.3.1 Basic or prototype emotions

Diverse writers have proposed that there are from two to twenty basic or prototype emotions (for example, [33], p. 8, [4], p. 10). The most common four appearing on these lists are: fear, anger, sadness, and joy. Plutchik [33] distinguished among eight basic emotions: fear, anger, sorrow, joy, disgust, acceptance, anticipation, and surprise.

Some authors have been concerned less with eight or so prototype emotions and refer primarily to dimensions of emotion, such as negative or positive emotions. Three dimensions show up most commonly. Although the precise names vary, the two most common categories for the 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.

Leidelmeijer [4] and Stein and Oatley [34] bring together evidence for and against the existence of basic emotions, especially universally. In a greater context, however, this problem of not being able to precisely define categories occurs all the time in pattern recognition and "fuzzy classification." Also, the universal issue is no deterrent, given the Negroponte speech recognition argument. It makes sense to simplify the possible categories of emotions for computers, so that they can start simply, recognizing the most obvious emotions.[10] The lack of consensus about the existence of precise or universal basic emotions does not interfere with the ideas I present below.

For affective computing, the recognition and modeling problems are simplified by the assumption of a small set of discrete emotions, or small number of dimensions. Those who prefer to think of emotions as continuous can consider these discrete categories as regions in a continuous space, or can adopt one of the dimensional frameworks. Either of these choices of representation comes with many tools, which I will say more about below.

Clynes's exclusivity principle of sentic states [14] suggests that we cannot express one emotion when we are feeling another – we cannot express anger when we are feeling hope. Clynes emphasized the "purity" of the basic sentic 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 also maintained that one can account for any emotion by a mixture of the principal emotions [33]. However, in the same article, Plutchik postulates that emotions are rarely perceived in a pure state. The distinctions between Plutchik and Clynes appear to be a matter of intensity and expression. One might argue that intensity is enhanced when one voluntarily *expresses* their sentic state, a conscious, cognitive act. As one strives for purity of expression, one moves closer to a pure state.

Given the human is in one sentic state, 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 [12] which might be observed in how you stand, walk, gesture, or otherwise behave.

---

[9]Perhaps this willingness may also be induced, ad infinitum.

[10]It is fitting that babies appear to have a less complicated repertoire of emotions than cogitating adults.
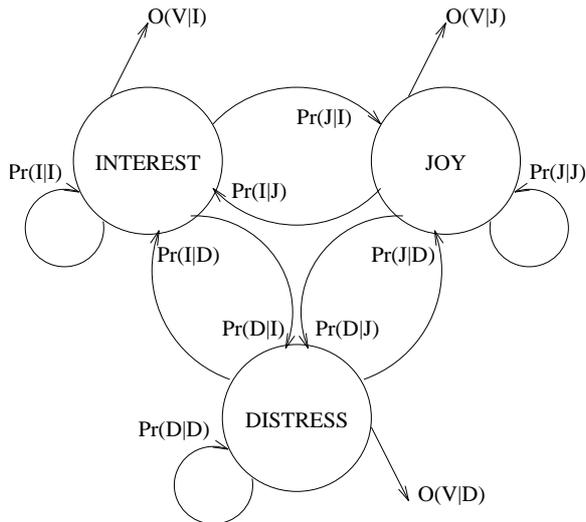
Figure 1: The state (here: Interest, Distress, or Joy) of a person cannot be observed directly, but observations which depend on a state can be made. The Hidden Markov Model shown here characterizes probabilities of transitions among three "hidden" states, (I,D,J), as well as probabilities of observations (measurable essentic forms, such as features of voice inflection, V) given a state. Given a series of observations over time, an algorithm such as Viterbi's [35] can be used to decide which sequence of states best explains the observations.

2.3.2 Affective state models

The discrete, hidden paradigm for sentic states suggests a number of possible models. Figure 1 shows an example of one such 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 HMM is trained on observations, which could be any measurements of sentic modulation. Different HMM's can be trained for different contexts or situations– hence, the probabilities and states may vary depending on whether you are at home or work, with kids, your boss, or by yourself. As mentioned above, the essentic form of certain motor system observations, such as finger pressure, voice, and perhaps even inspiration and expiration during breathing, will vary as a function of the different states. Its dynamics are the observations used for training and for recognition. Since an HMM can be trained on the data at hand for an individual or category of individuals, the problem of universal categories does not arise.

HMM's can also be adapted to represent mixture emotions. Such mixtures may also be modeled (and tailored to an individual, and to context) by the cluster-based probability model of Popat and Picard [36]. In such a case, high-dimensional probability distributions could be learned for predicting sentic states or their mixtures based on 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. Artificial neural nets, and explicit mixture models may also be useful

for sentic state modeling. Camras [37] has also 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.

If explicit dimensions are desired, one could compute eigenspaces of features of the observations and look for identifiable "eigenmoods." These might correspond to either pure or mixture emotions. The resulting eigenspaces would be interesting to compare to the spaces found in factor analysis studies of emotion; such spaces usually include the axes of valence (positive, negative), and arousal (calm, excited). 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 be characterized in these spaces, capturing the dynamic aspects of emotion.

Given one of these state-space or dimension-space models, trained on individual motor outputs, then features of unknown motor outputs can be collected in time, and used with tools such as maximum a posterior decision-making to recognize a new unclassified emotion. Because the recognition of sentic state can be set up as a supervised classification problem, i.e. one where classes are specified a priori, a variety of pattern recognition and learning techniques are available [38], [39].

2.3.3 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 [40] 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 [40]:

> 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.

There are many types of HMM's, including the recent Partially Observable Markov Decision Processes (POMDP's) which give a "reward" associated with executing a particular action in a given state [41], [42], [43]. These models also permit observations at each state which are actions [44], and hence could incorporate not only autonomic measures, but also behavioral observations.

There have also been computational models of emotion proposed that do not involve emotion recognition per se, but rather aim to mimic the mechanisms by which emotions might be produced, for use in artificial intelligence systems (See Pfeifer [45] for a nice overview.) The artificial intelligence emphasis has been on modeling emotion *in* the computer (what causes the computer's emotional state to change, etc.); let's call these "synthesis" models. In contrast, my emphasis is on equipping the computer to express and recognize affective state. Expres-

sion is a mode-specific (voice, face, etc.) synthesis problem, but recognition is an analysis problem. Although models for synthesis might make good models for recognition, especially for inferring emotions which arise after cognitive evaluation, but which are not strongly expressed, nonetheless, they can often be unnecessarily complicated. Humans appear to combine both analysis and synthesis in recognizing affect. For example, the affective computer might hear the winning lottery number, know that it's your favorite number and you played it this week, and predict that when you walk in, you will be elated. If you walk in looking distraught (perhaps because you can't find your ticket) the emotion synthesis model would have to be corrected by the analysis model.

3 Things Better Left Unexplored?

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

A curious student posed to me the all-too-infrequently asked important question, if affective computing isn't a topic "better left unexplored by humankind." At the time, I was thinking of affective computing as the two cases of computers being able to recognize emotion, and to induce emotion. My response to the question was that there is nothing wrong with either of these; the worst that could come of it might be emotion manipulation for malicious purposes. Since emotion manipulation for both good and bad purposes is already commonplace (cinema, music, marketing, politics, etc.), why not use computers to better understand it?" However these two categories of affective computing are much less sinister than the one that follows.

3.1 Computers that kill

The 1968 science fiction classic movie "2001: A Space Odyssey" by Kubrick and Clarke, and subsequent novel by Clarke, casts a different light on this question. A HAL 9000 computer, "born" January 12, 1993, is the brain and central nervous system of the spaceship Discovery. The computer, who prefers to be called "Hal," has verbal and visual capabilities which exceed those of a human. Hal is a true "thinking machine," in the sense of mimicking both cortical and limbic functions, as evinced by the 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 don't think anyone can truly answer."

As the movie unfolds, it becomes clear that Hal is not only articulate, but capable of both expressing and perceiving feelings, by his lines:

> "I feel much better now."

> "Look, Dave, I can see you're really upset about this."

But Hal goes beyond expressing and perceiving feelings. In the movie, Hal becomes afraid of being disconnected. 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, Dave Bowman, with whom Hal said earlier he "enjoys a stimulating relationship." Hal not only

passes the Turing test, he has the ability to kill the person administering it.

The theme bears repeating. When everyone is complaining that computers can't think like humans, an erratic genius, Dr. Daystrom, comes to the rescue in the episode, "The Ultimate Computer," from *Star Trek, The Original Series*. Daystrom impresses his personal engrams on the highly advanced M5 computer, which is then put in charge of running the ship. It does such a fine job that soon the crew is making Captain Kirk uncomfortable with jokes about how he is no longer necessary.

But, soon the M5 also concludes that people are trying to "kill" it; and this poses an ethical dilemma. Soon it has Kirk's ship firing illegally at other Federation ships. Desperate to convince the M5 to return the ship to him before they are all destroyed, Kirk tries to convince the M5 that it has committed murder and deserves the death penalty. The M5, with Daystrom's personality, reacts first with arrogance and then with remorse, finally surrendering the ship to Kirk.

3.1.1 A dilemma

The message is serious: a computer that can express itself emotionally will some day act emotionally, and the consequences will be tragic. Objection to such an "emotional computer" based on fear of the consequences parallels the "Heads in the Sand" objection, one of nine objections playfully proposed and refuted by Turing in 1950 to the question "Can machines think?"[46].

This objection leads, perhaps more importantly, to a dilemma that can be stated more clearly when one acknowledges the role of the limbic system 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. [5]

In fact, it is commonly acknowledged that machine intelligence cannot be achieved without an ability to understand and express affect. Negroponte, referring to Hal, says

> HAL had a perfect command of speech (understanding and enunciating), excellent vision, and humor, which is the supreme test of intelligence. [10]

But with this flexibility, intelligence, and creativity, comes unpredictability. Unpredictability in a computer, and the unknown in general, evoke in us a mixture of curiosity and fear.

Asimov, in "The Bicentennial Man" [47] presents three laws of behavior for robots, ostensibly to solve this dilemma and prevent the robots from bringing harm to anyone. 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. The law-based robot would be severely handicapped in its decision-making ability, not too differently from the frontal-lobe patients of Damasio.

Can we create computers that will recognize and express affect, exhibit humor and creativity, and never bring about harm by emotional actions?

3.2 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 [48]

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.

Although expressing and recognizing affect are important for computer-human interaction, building emotion into the motivational behavior of the computer is a different issue. "Emotional" when it refers to people or to computers, usually connotes an undesirable reduction in sensibilities. Interestingly, in the popular series *Star Trek, The Next Generation*, the affable android Data was not given emotions, although he is 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. Both Data and his brother appear to have the ability to kill, but Data cannot kill out of malice.

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 knows. Or, computers with the power to kill should not have emotions,[11] 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. Clearly, computers would benefit from development of ethics, morals, and perhaps also of religion.[12] These developments are important even without the amplifier of affect.

3.2.1 Four cases

What are the cases that arise in affective computing, and how might we proceed, given the scenarios above? Table 1 presents four cases, but there are more. For example, I omitted the rows "Computer can/can't 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't act based on emotions" for the reasons described above. The ethical and philosophical implications of such "emotionally-based computers" take us beyond the scope of this essay; these possibilities are not included in Table 1 or addressed in the applications below.

This leaves the four cases in Table 1, which I'll describe briefly:

---

[11]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." Do you want to wait for your computer to *feel interested* before it will listen to you? I doubt electronic computers will *have feelings*, but I recognize the parallels in this statement to debates about machines having consciousness. If I were a betting person I would wager that someday computers will be better than humans at feigning emotions. But, should a computer be allowed emotional autonomy if it could develop a bad attitude, put its self-preservation above human life, and bring harm to people who might seek to modify or discontinue it? Do humans have the ability to endow computers with such "free will?" Are we ready for the computer-rights activists? These questions lead outside this essay.

[12]Should they fear only their maker's maker?

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, enabling it to adjust its response in ways that might, for example, make it a better teacher and more useful assistant. It 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.

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.

3.3 Affective symmetry

In crude videophone experiments we wired up at Bell Labs a decade ago, my colleagues and I learned that we 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.

Computers that read emotions (i.e., infer hidden emotional state based on physiological and behavioral observations) should also show us what they are reading. In other words, affective interaction with computers can easily give us 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.

If computer scientists persevere in giving computers internal emotional states, then I suggest that these states should be observable. If Hal's sentic 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 2001 storyline would not work.

Will we someday hurt the feelings of our computers? I hope this won't be possible; but, if it is a possibility, then the computer should not be able to hide its feelings from us.

4 Affective Media

Below I suggest scenarios for applications of affective computing, focusing on the three cases in Table 1 where computers can perceive and/or express affect. The scenarios below assume modest success in correlating observations of an individual with at least a few appropriate affective states.

4.1 Entertainment

One of the world's most popular forms of entertainment is large sporting events – especially World Series, World Cup, Super Bowl, etc. One of the pleasures that people receive from these events (whether or not their team wins) is the opportunity to express intense emotion as part of a large crowd. A stadium is one of the only places where an adult can jump up and down cheering and screaming, and be looked upon with approval, nay accompanied. Emotions, whether or not acknowledged, are an essential part of being entertained.

> Do you feel like I do? – Peter Frampton

Although I'm not a fan of Peter Frampton's music, I'm moved by the tremendous response of the crowd in his live recorded performance where he asks this question repeatedly, with increasingly modified voice. Each time he poses the question, the crowd's excitement grows. Are they mindless fans who would respond the same to a mechanical repeating of the question, or to a rewording of "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 waggling 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 dog, or other character it changed into) vs. the predictable errors the submarine would make when we did not steer correctly? What makes one experience pleasurably more engaging than another?

Clynes's "self-generating principle" indicates that the intensity of a sentic state is increased, within limits, by the repeated, arrhythmic generation of essentic form. Clynes has carried this principle forward and developed a process of "sentic cycles" whereby people (in a controlled 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.

These effects occur also in role-playing, 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. Entertainment can serve to expand our emotional dynamic range.

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, "pulls one in" so to speak, becomes apparent in your face and posture. It need not draw forth a roar, or a lot of waggling of reflectors, or a lot

of pushing of buttons as in the interactive theatres coming soon from Sony. Nonetheless, the performer is keen at sensing how the audience is responding, and is, in turn, affected by their response.

Suppose that audience response could be captured 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. The floors in the intermission area might be live compositions, waiting to sense the mood of the audience and amplify it with music. New musical instruments, such as Tod Machover's hyperinstruments, might also be equipped to sense affect directly, augmenting the modes of expression available to the performer.

4.2 Expression

> *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 [14]

Clynes [14] 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 (2) then 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. [14]

Clynes attributes the power of Casal's performance to the purity and preciseness of the essentic form. Faithfulness to the purest inner form produces beautiful results.

With sentic recognition, the computer music teacher could not only keep you interested longer, but it could also give feedback as you develop preciseness of expression. Through measuring essentic form, perhaps through finger pressure, foot pressure, or measures of inspiration and expiration as you breathe, it could help you compare aspects of your performance that have never been measured or understood before.

Recently, Clynes [29], 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 what I call a "Musical Turing test" [13] 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

---

[13] 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.

ranked the computer's as second or third on average. Clynes'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 the finest language of emotion, music.

4.2.1 Expressive mail

Although sentic states may be subtle in their modulation of expression, they are not subtle in their power to communicate, and correspondingly, to persuade. When sentic (emotion) modulation is missing, misunderstandings are apt to occur. Perhaps nothing has brought home this point more than the tremendous reliance of many people on Email (electronic mail) 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, namely smileys such as :-) and ;-((turn your head to the left to recognize). Even so, these icons are very limited, and Email communication consequently carries much less information than a phone call.

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 such as skin conductivity, temperature, and pulse, all observations which may be combined to identify emotional state.

Encoding sentic state instead of the specific mode of expression permits the message to be transmitted to the widest variety of receivers. Regardless whether the receiver handled visual, auditory, text, or other modes, it could transcode the sentic message appropriately.

4.3 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

It is the job of the director to create onstage or onscreen, a mood, a precise essentic form, that provokes a desired affect in the audience. "Method acting" inspired by the work of Stanislavsky, is based on the recognition that the actor that feigns an emotion is not as convincing as the actor that is filled with the emotion; the latter has a way of engaging you vicariously in the character's emotion, provided you are willing to participate. Although precisely what constitutes an essentic form, and precisely what provokes a complementary sentic state in another is hard to measure, there is undoubtably a power we have to transfer genuine emotion. We sometimes say emotions are contagious. Clynes suggests that the purer the underlying essentic form, the more powerful its ability to persuade.

4.3.1 Skip ahead to the interesting part

My research considers how to help computers "see" like people see, with all the unknown and complicated aspects human perception entails. One of the 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 "skip ahead to the

dinosaurs." A longer term, and much harder goal, is to "make a long story short." How does one teach a computer to 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 harder challenge.

We have recently built some of the first tools that enable computers to assist humans in annotating video, attaching descriptions to images that the person and computer look at together [49]. 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.[14] Affective computing can help systems such as this begin to learn not only which content is most interesting, but what emotions might be stirred by the content.

Interest is related to 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) [50].

In fact, unlike trying to search for a shot that has a particular verbal description of content (where choice of what is described may vary tremendously and be quite lengthy), affective annotations, especially in terms of a few basic emotions or a few dimensions of emotion, could provide a relatively compact 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!." It is not uncommon that someone might want to retrieve the thrilling shots. Affective annotations would be verbal descriptions of these primarily nonverbal events.

Instead of annotating, "this is a sunny daytime shot of a student getting his diploma and jumping off the stage" we might annotate "this shot of a student getting his diploma makes people grin." Although it is a joyful shot for most, however, it may not provoke a grin for everyone (for example, the mother whose son would have been at that graduation if he were not killed the week before.) Although affective annotations, like content annotations, will not be universal, they will still help reduce time searching for the "right scene." Both types of annotation are potentially powerful; we should be exploring them in digital audio and visual libraries.

4.4 Environments

Sometimes you like a change of environment; sometimes it drives you crazy. These responses apply to all environments – not just your building, home, or office, but also your computer 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 [51] 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 and cognitive computing have been largely concerned with measuring information in the first and second categories. Affective computing addresses the third.

---

[14]Computers have a hard time learning similarity, so this system tries to adapt to a user's ideas of simility - whether perceptual, semantic, or otherwise.

Stewart Brand's book "Buildings that Learn" [52] 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, because these are written or verbal reports, and the language of feelings is so inexact, these surveys are limited in their ability to capture what is really felt. Brand notes that surveys also lose in the sense that they occur at a much later time than the actual experience, and hence may not recall what was actually felt.

In contrast, measuring sentic responses of newcomers to the building could tell you how the customers feel when they walk in your bank vs. into the competitor's bank. Imagine surveys where newcomers are asked to express their feelings when they enter your building for the first time, and an affective computer records their response.

After being in a building awhile, your feelings in that space are no longer likely to be influenced by its structure, as that 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 match or help create an appropriate mood. In your car, your digital disc jockey could play for you the music that you'd find most agreeable, depending on your mood at the end of the day. "Computer, please adjust for a peaceful ambience at our party tonight."

4.5   Aesthetic pleasure

> 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 [53]

> Art does not think logically, or formulate a logic of behaviour; it expresses its own postulate of faith. If 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 [54]

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. This issue is deeply complex and elusive; this is true, in part, because affect plays a primary role.

For scientists and others who do not see why a computer needs to be concerned with aesthetics, consider a scenario where a computer is assembling a presentation for you. The computer will be able to search digital libraries all over the world, looking for images and video clips with the content you request. 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 choose a set of "good" ones to show you. How do you teach it what is "good?" Can something be measured intrinsically in a picture, sculpture, building, piece of music, or flower arrangement that will indicate its beauty and appeal?

4.5.1   Hidden forms

Clynes identified visual traces of essentic form 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 the measured finger-pressure forms, are indicative of the 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.

Many today refute the view that we could find some combination of primitive elements in a picture that corresponds to an emotion – as soon as we've found what we think is the form in one image, we imagine we can yank it into our image-manipulation software and construct a new image around it, one that does not communicate the same emotion. Or we can search the visual databases of the net and find visually similar patterns to see if they communicate the same emotion. It seems it would be easy to find conflicting examples, especially across cultural, educational, and social strata. However, to my knowledge, no such investigation has been attempted yet. Moreover, finding ambiguity during such an investigation would still not imply that 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 see the lines she saw, then presumably it could help her find new music she would like. For certain synesthetes, rules might be discovered to predict these aesthetic feelings. What about for the rest of us?

The idea from Cytowic's studies is that perception, in all of us, passes through the limbic system. However, synesthetes are also aware of the perceived form while it is passing through this intermediate stage. Perhaps the limbic system is where Clynes's hypothesized "essentic form" resides. Measurements of human essentic forms may provide the potential for "objective" recognition of the aesthete. Just as lines going up and at an angle means the woman likes the music, 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 place 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 another scenario.

4.5.2   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 gag – "how could anybody like any of these?"

People's preferences and tastes for what they like 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 department store knows that everyone will not find the same garment equally beautiful. A universal predictor of what everyone would like is absurd.

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

expressive intent. 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 object – garment or artwork, that influences our judgement. But what functions of appearance should be measured? Perhaps if the lines in that print were straighter, it would be too boring for you. But you might treasure the bold lines on that piece of furniture.

There are a lot of problems with trying to find something to measure, be it in a sculpted style, painting, print, fabric, or room decor. Ordinary pixels and lines don't induce aesthetic feelings on their own, unless, perhaps it is a line of Klee, used to create an entire figure. Philosophers such as Susanne K. 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. [55]

But, neither do we experience aesthetic feelings and aesthetic pleasure *without* the pixels and lines and notes and rhythms. And, 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. Just thinking of a magnificent painting or piece of music does not usually arouse the same feelings as when one is actually experiencing the work, but it may arouse similar, fainter feelings. Indeed, Beethoven 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 arousal.

To give computers personal recognition of what we think is beautiful will probably require lots of examples of things that we do and don't like, and the ability of the computer to incorporate affective feedback from us. The computer will need to explore features that detect similarities among those examples that we like[15], and distinguish these from features common to the negative examples. Ultimately, it can cruise the network catalogs at night, helping us shop for clothes, furniture, wallpaper, music, gifts, and more.

Imagine a video feedback system that adjusts its image or content until the user is pleased. Or a computer director or writer that adjusts the characters in the movie, until the user empathically experiences the story's events. Such a system might also identify the difference between the user's sadness due to story content, e.g., the death of Bambi's mom, and the user's unhappiness due to other factors – possibly a degraded color channel, or garbled soundtrack. If the system were wearable, and perhaps seeing everything you see [56], then it might correlate visual experiences with heart rate, respiration, and other forms of sentic modulation. Affective computing will play a key role in better aesthetic understanding.

4.5.3 Design

> You can't 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

---

[15]Perhaps by having a "society of models" that looks for similarities and differences, as in[49].

Have you asked a designer how they arrived at the final design? Of course, there are design principles and constraints on function that influenced 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" that fires when they've got it right.

Although we can measure qualities of objects, the space between them, and 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 (which is usually much larger than four dimensions). 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. But the key missing objective of these systems is the goal of arousing the user – arousing to provoke attention, interest, and memory.

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, expressing their feelings about the computer's design suggestions.)

Aesthetic success is communicated via feelings. 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, and this brings you joy. Or a design solves a problem you have and now you feel relief. Ideally, in the end it brings you to a new state that feels better than the one you were in.

Although the computer does not presently know how to lead a designer to this satisfied state, there's no reason it couldn't 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. Frequently, it is precisely the sentic response that is targeted during design – we ought to equip computers with the ability to recognize this.

> 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 that perceive what you perceive, and recognize personal responses. In the most personal form, these are computers that could accompany you at all times.

4.6 Affective wearable computers

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 now available in Harvard Square. Of course these armpit heat-to-color transformers don't 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 in your belt, a miniature videocamera and holographic display on your eyeglasses, and

more. Wearables may fulfill some of the dreams espoused by Clynes when he coined the word "cyborg" [57]. Wearable computers can augment your memory (any computer accessible information available as you need it) [58] 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 their name and where you last met. Signals can be passed from one wearable to the other through your conductive "BodyNet," [59]. A handshake could instantly pass to my online memory the information on your business card.[16]

One of my favorite science fiction novels [60] features a sentient being named Jane that speaks from a jewel in the ear of Ender, the hero of the novel. 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.

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.

The bandwidth we have for communicating thoughts and feelings to humans should 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, in an effort to better understand you. You could choose whether or not your wearable would reveal these personal clues of your emotional state to anyone.

> 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 that the local florist just received a delivery of your spouse's favorite protea. A mood detector might even make suggestions about what foods to eat, so called "mood foods" [61].

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 that senses your posture might gently remind you to correct a bad habit after back surgery. Wearables that measure other physiological responses can help you identify causes of stress and anxiety, and how well your body is responding to these.[17] 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 private reflection – "I sense more joy in you tonight."

You have heard someone with nary a twitch of a smile say "pleased to meet you" and perhaps wondered about their sincerity. Now imagine you are both donned with wearable affective computers that you allow to try to recognize your emotional state. Moreover, suppose we permitted these wearables to communicate between people, and whisper in your ear, "he's not entirely truthful." Not only would we quickly find out how reliable polygraphs are, which usually measure heart rate, breathing rate, and galvanic skin response, but imagine the implications for communication (and, perhaps, politics).

With willful participants, and successful affective computing, the possibilities are only limited 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 take careful notes for you, assuming that your mind is wandering. Games might add points for courage. Your wearable might coax during a workout, "keep going, I sense healthy anger reduction." Wearables that were allowed to network might help people reach out to contact those who want to be contacted those of unbearable loneliness, the young and the old [29].

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. But why not let the lilt in your walk on the way to your car (perhaps sensed by affective sneakers) tell the digital disc jockey to put on happy tunes?

4.6.1 Implications for emotion theory

Despite a number of significant works, emotion theory is far from complete. In some ways, one might even say it is stuck. 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 would be real life observation, recently believed to be impossible [17].

However, a computer you wear, that attends to you during your waking hours, could notice what you eat, what you do, what you look at, and what emotions you express. Computers excel at amassing and, to some extent, analyzing information and looking for consistent patterns. Although ultimate interpretation and use of the information should be left to the wearer and to those in whom the wearer confides, this information could provide tremendous sources to researchers interested in human diet, exercise, activity, and mental health.

5 Summary

Emotions have a major impact on essential cognitive processes; neurological evidence indicates they are not a luxury. I have highlighted several results from the neurological literature which indicate that emotions play a necessary role not only in human creativity and intelligence, but also in rational human thinking and decision-making. Computers that will interact naturally and intelligently with humans need the ability to at least recognize and express affect.

Affective computing is a new field, with recent results primarily in the recognition and synthesis of facial expression, and the synthesis of voice inflection. However, these modes are just the

---

[16]It could also pass along a virus, but technology has had success fighting computer viruses, in contrast with the biological variety.

[17]See [62] for a discussion of emotions and stress.

tip of the iceberg; a variety of physiological measurements are available which would yield clues to one's hidden affective state. I have proposed some possible models for the state identification, treating affect recognition as a dynamic pattern recognition problem.

Given modest success recognizing affect, numerous new applications are possible. Affect plays a key role in understanding phenomena such as attention, memory, and aesthetics. I described areas in learning, information retrieval, communications, entertainment, design, health, and human interaction where affective computing may be applied. In particular, with wearable computers that perceive context and environment (e.g. you just learned the stock market plunged) as well as physiological information, there is the promise of gathering powerful data for advancing results in cognitive and emotion theory, as well as 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. Based on this evidence, to build computers that make intelligent decisions may require building computers that "have emotions."

I have 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, the maker, may be eliminated by our creation. I have argued for 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 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, feelings, and freedom.

Acknowledgments

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] C. E. Izard, "Four systems for emotion activation: Cognitive and noncognitive processes," *Psychological Review*, vol. 100, no. 1, pp. 68–90, 1993.

[4] K. Leidelmeijer, *Emotions: An experimental approach*. Tilburg University Press, 1991.

[5] R. E. Cytowic, *The Neurological Side of Neuropsychology*. Cambridge, MA: MIT Press, 1995. To Appear.

[6] A. R. Damasio, *Descartes' Error: Emotion, Reason, and the Human Brain*. New York, NY: Gosset/Putnam Press, 1994.

[7] P. N. Johnson-Laird and E. Shafir, "The interaction between reasoning and decision making: an introduction," *Cognition*, vol. 49, pp. 1–9, 1993.

[8] 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.

[9] 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.

[10] N. Negroponte, *Being Digital*. New York: Alfred A. Knopf, 1995.

[11] E. D. Scheirer, 1994. Personal Communication.

[12] M. Lewis, "Self-conscious emotions," *American Scientist*, vol. 83, pp. 68–78, Jan.-Feb. 1995.

[13] B. W. Kort, 1995. Personal Communication.

[14] D. M. Clynes, *Sentics: The Touch of the Emotions*. Anchor Press/Doubleday, 1977.

[15] 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.

[16] R. Plutchik and H. Kellerman, eds., *Emotion Theory, Research, and Experience*, vol. 1–5. Academic Press, 1980–1990. Series of selected papers.

[17] 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.

[18] R. B. Zajonc, "On the primacy of affect," *American Psychologist*, vol. 39, pp. 117–123, Feb. 1984.

[19] G. Mandler, "The generation of emotion: A psychological theory," in *Emotion Theory, Research, and Experience* (R. Plutchik and H. Kellerman, eds.), vol. 1, Theories of Emotion, Academic Press, 1980.

[20] G. D. de Boulogne, *The Mechanism of Human Facial Expression*. New York, NY: Cambridge University Press, 1990. Reprinting of original 1862 dissertation.

[21] P. Ekman and W. Friesen, *Facial Action Coding System*. Consulting Psychologists Press, 1977.

[22] I. A. Essa, *Analysis, Interpretation and Synthesis of Facial Expressions*. PhD thesis, MIT Media Lab, Cambridge, MA, Feb. 1995.

[23] 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.

[24] 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.

[25] R. W. Picard, "Content access for image/video coding: "The Fourth Criterion"," Tech. Rep. 295, MIT Media Lab, Perceptual Computing, Cambridge, MA, 1994.

[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] J. E. Cahn, "The generation of affect in synthesized speech," *Journal of the American Voice I/O Society*, vol. 8, July 1990.

[28] 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.

[29] M. Clynes, 1995. Personal Communication.

[30] R. E. Cytowic, 1994. Personal Communication.

[31] M. Minsky, *The Society of Mind*. New York, NY: Simon & Schuster, 1985.

[32] B. Burling, 1995. Personal Communication.

[33] 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.

[34] 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.

[35] L. R. Rabiner and B. H. Juang, "An introduction to hidden Markov models," *IEEE ASSP Magazine*, pp. 4–16, Jan. 1986.

[36] K. Popat and R. W. Picard, "Novel cluster probability models for texture synthesis, classification, and compression," in *Proc. SPIE Visual Communication and Image Proc.*, vol. 2094, (Boston), pp. 756–768, Nov. 1993.

[37] L. A. Camras, "Expressive development and basic emotions," *Cognition and Emotion*, vol. 6, no. 3 and 4, 1992.

[38] R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*. Wiley-Interscience, 1973.

[39] C. W. Therrien, *Decision Estimation and Classification*. New York: John Wiley and Sons, Inc., 1989.

[40] 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.

[41] 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.

[42] W. S. Lovejoy, "A survey of algorithmic methods for partially observed Markov decision processes," *Annals of Operations Research*, vol. 28, pp. 47–66, 1991.

[43] 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.

[44] T. Darrell, "Interactive vision using hidden state decision processes." PhD Thesis Proposal, 1995.

[45] 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.

[46] A. M. Turing, "Computing machinery and intelligence," *Mind*, vol. LIX, pp. 433–460, October 1950.

[47] I. Asimov, *The Bicentennial Man and Other Stories*. Garden City, NY: Doubleday Science Fiction, 1976.

[48] 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.

[49] R. W. Picard and T. M. Minka, "Vision texture for annotation," *Journal of Multimedia Systems*, vol. 3, pp. 3–14, 1995.

[50] B. Reeves, "Personal communication, MIT Media Lab Colloquium," 1995.

[51] 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.

[52] S. Brand, *How buildings learn: what happens after they're built*. New York, NY: Viking Press, 1994.

[53] 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'.

[54] A. Tarkovsky, *Sculpting in Time: Reflections on the Cinema*. London: Faber and Faber, 1989. ed. by K. Hunter-Blair.

[55] S. K. Langer, *Mind: An Essay on Human Feeling*, vol. 1. Baltimore: The Johns Hopkins Press, 1967.

[56] S. Mann, "'See the world through my eyes,' a wearable wireless camera," 1995. http://www-white.media.mit.edu/~steve/netcam.html.

[57] M. Clynes and N. S. Kline, "Cyborgs and space," *Astronautics*, vol. 14, pp. 26–27, Sept. 1960.

[58] T. E. Starner, "Wearable computing," Perceptual Computing Group, Media Lab 318, MIT, Cambridge, MA, 1995.

[59] N. Gershenfeld and M. Hawley, 1994. Personal Communication.

[60] O. S. Card, *Speaker for the Dead*. New York, NY: Tom Doherty Associates, Inc., 1986.

[61] J. J. Wurtman, *Managing your mind and mood through food*. New York: Rawson Associates, 1986.

[62] G. Mandler, *Mind and Body: Psychology of Emotion and Stress*. New York, NY: W. W. Norton & Company, 1984.

16