

How emotion is made and measured

Kirsten Boehner^{a,*}, Rogério DePaula^b, Paul Dourish^c, Phoebe Sengers^a

^a*Cornell Information Science, Cornell University, 301 College Avenue, Ithaca NY 14850, USA*

^b*Intel Research, Av. Dr. Chucuri Zaidan 940, 10 floor, São Paulo 04583-110, Brazil*

^c*Department of Informatics, UC Irvine, Irvine CA 92697-3440, USA*

Available online 19 January 2007

Abstract

How we design and evaluate for emotions depends crucially on what we take emotions to be. In affective computing, affect is often taken to be another kind of information—discrete units or states internal to an individual that can be transmitted in a loss-free manner from people to computational systems and back. While affective computing explicitly challenges the primacy of rationality in cognitivist accounts of human activity, at a deeper level it often relies on and reproduces the same information-processing model of cognition. Drawing on cultural, social, and interactional critiques of cognition which have arisen in human–computer interaction (HCI), as well as anthropological and historical accounts of emotion, we explore an alternative perspective on emotion as interaction: dynamic, culturally mediated, and socially constructed and experienced. We demonstrate how this model leads to new goals for affective systems—instead of sensing and transmitting emotion, systems should support human users in understanding, interpreting, and experiencing emotion in its full complexity and ambiguity. In developing from emotion as objective, externally measurable unit to emotion as experience, evaluation, too, alters focus from externally tracking the circulation of emotional information to co-interpreting emotions as they are made in interaction.

© 2006 Elsevier Ltd. All rights reserved.

Keywords: Affective computing; Affective evaluation

1. Introduction

A social, interactional approach to understanding cognition in human–computer interaction (HCI) has emerged in the last 20 years in contrast to the dominant information-processing approach to capturing, modeling, augmenting, and supporting human activity. The recent emphasis on the importance of emotion for cognition further advances these arguments to look “beyond the cognitive” and to understand new aspects of human experience. Nevertheless, there is a critical difference between the turn to emotions and the turn to the social in HCI. While the social and cultural approaches attempt to deconstruct conventional approaches to cognition (and in particular the underlying cognitivist computational

claim on mind), the recent exploration of the role of emotions leaves traditional cognitivism intact, and in fact depends on it as the base for adding emotional understandings.

Emotion, in the informational model, is a dual of cognition, but it is nonetheless the same sort of phenomenon—an internal, individual, and delineable phenomenon, which operates in concert with and in the context of traditional cognitive behavior. That is, while emotion is proposed as a supplement to traditional cognitive accounts, it is often nonetheless located within the same information-processing frame. For this reason, emerging understandings of emotion can suffer from similar problems as purely cognitive approaches in the past—that is, difficulty in accounting for and adequately incorporating an understanding of everyday action as situated in social and cultural contexts that give them meaning.

Rather than relying primarily on an informational account, we offer and critically explore an interactional account of emotion and the role that it plays in action and

*Corresponding author. Tel.: +1 607 254 5396; fax: +1 607 255 5196.

E-mail addresses: kab18@cornell.edu (K. Boehner), rogerio.depaula@intel.com (R. DePaula), jpd@ics.uci.edu (P. Dourish), sengers@cs.cornell.edu (P. Sengers).

practice. As argued by Boellstorff and Lindquist (2004), citing Rosaldo (1984), “feelings are not substances to be discovered in our blood but social practices organized by stories that we both enact and tell.” The production and interpretation of emotion—of national pride, justifiable anger, or shame—is social and cultural in origin. We take emotion as a social and cultural product experienced through our interactions.

This interactional approach to emotions extends current HCI agendas, and in particular current affective computing research, in three ways. First, this approach sees emotions as culturally grounded, dynamically experienced, and to some degree constructed in action and interaction. This expands the ontological view of emotions as informational units that are internally constructed and subsequently delivered. Second, as an interface paradigm, an interactional approach moves the focus from helping computers to better understand human emotion to helping people to understand and experience their own emotions—the raw elements and perceptions of emotions, the constructed conceptions of these emotions, and the resulting effects such as behavioral or cognitive changes. Finally, the interactional approach leads to new design and evaluation strategies for devices. Systems inspired by the interactional approach to emotion emphasize the expression of emotion in a co-constructed, co-interpreted fashion. Measures of success for such systems therefore do not focus on whether the systems themselves deduce the ‘right’ emotion but whether the systems encourage awareness of and reflection on emotions in users individually and collectively.

In this paper, we consider the turn in HCI to affective computing and, in particular, the different expectations, commitments, and entailments that occur in moving from an informational to an interactional model. We suggest that the informational model is rooted in a longer laboratory-science tradition of studying emotion in which subjective experience is considered suspect, to be replaced by objective measures, but that this tradition does not address salient aspects of emotion as experienced in interaction. In moving out of this tradition, we show how this shift necessarily entails a new understanding of evaluation. Affect as interaction does not simply suggest new methods for evaluation; it also calls into question some of the epistemological basis upon which evaluation traditionally sits.

2. Emotions, affective computing, and HCI

In this section, we describe the model of emotion that underlies current affective computing. We trace its heritage in laboratory studies of emotion, and show how the model of emotion as biological, measurable, and objectively present is taken up and extended in the context of information processing to become emotion-as-information. We will argue that the laboratory-science heritage that informs many affective computing approaches may unintentionally sideline precisely the issue at the heart of

affective computing—how people interactively construct and subjectively experience their emotions. Drawing on critiques of the information-processing model of cognition in HCI and communication, as well as on the anthropology and history of emotions, we develop the alternative model of affect as interaction.

2.1. Affect as a measurable, biological fact

Unsurprisingly, HCI is not the first research area to grapple with the problem of emotion as an aspect of human behavior. Similarly, it is not the first research area to have to find a way to make emotion experimentally available and conceptually compatible with existing research models. In order to understand how we in HCI conceive of what an emotion is, we first need to understand the roots of emotion in affective computing’s dominant heritage, the laboratory study of emotion.

Emotion presents a particular challenge for scientific study because, in modern Western cultural traditions, emotion is generally defined in contrast to cognition—and by extension science itself. As Otniel Dror (1999a, b, 2001) argues, the science of affect has been strongly informed by modern Western notions of what an emotion is and, more importantly, is *not*—i.e. rational—and how it must be reshaped to become amenable to scientific study. In Dror’s (2001) account, affect became of interest to scientists in the late 19th century, but its adoption as an object of study was hampered by the cultural status of emotion. Emotion was considered feminine, spiritual, and out of control, while the scientific laboratory was a masculine, physically grounded and highly controlled space. The notion of science as a rational activity therefore clashed with the study of difficult-to-grasp, fuzzy, uncontrolled emotions.

Similar discomfort with the potentially unruly, feminine, and non-rational nature of emotion can be seen in contemporary affective computing. In the preface to her landmark work that launched the field, Picard unpacks her own discomfort with emotions: “I never expected to write a book addressing emotions. My education has been dominated by science and engineering, and based on axioms, laws, equations, rational thinking, and a pride that shuns the ‘touchy-feely.’ Being a woman in a field containing mostly men has provided extra incentive to cast off the stereotype of ‘emotional female’ in favor of the logical behavior of a scholar” (Picard, 1997, p. x). Picard’s introduction does substantial work to defend against the idea that emotion must be negatively valenced—out of control, socially inappropriate, anti-rational, and anti-scientific. As we will unpack in more detail below, Picard rehabilitates emotion for computing by connecting it to rational cognition, thereby making it recognizably scientific.

As Dror (2001) documents, a similar move in the 19th century to reconcile emotion with science and, by extension, rational cognition, led to a fundamental shift in the notion of what emotion is. While ‘natural’ emotions might be ill-defined and out-of-control, in the laboratory

emotions could be translated through physiological measurement into numbers and graphs. These objective, distanced representations, Dror argues, were safe for study in a way that subjective, personal emotional experiences were not: “It empowered experimenters who wished to study, provoke, and release dangerous emotions inside the laboratory—without corrupting or disrupting knowledge production. It created an emotion that did not threaten the laboratory’s self-representation as an emotion-free space” (p. 359). The notion of emotion as physiologically measurable and mathematically documentable in turn led to the idea that emotion is fundamentally a natural, biological fact, something objectively observable, definable, and containable.

The insight that emotion as a scientific object of study—as a biological fact, individually experienced, quantitatively measurable—rests on our own cultural notions of ‘emotion’ is not in and of itself particularly problematic. After all, science is, like all human activity, part of our culture, and this does not prevent it from giving us useful knowledge. What *is* problematic for researchers interested in documenting and supporting complex human experiences of emotion, though, is the way in which during the process of becoming ‘studiable’ the definition of emotion has been altered to fit a particular conception of what science ought to be: rational, well-defined, and culturally universal. In the process, those aspects of emotional experience that do not fit this view of science—that which is not objectively accessible and measurable, that which is personal and idiosyncratic, that which varies over cultures and over time—are left out of discussion. To put it another way, emotion is not thought of as biological, measurable, and objectively present because scientists found it to exist in the world that way, but because 19th-century scientists could not imagine studying it scientifically any other way. Nevertheless, we will argue below that this model of emotion remains central to many contemporary affective computing approaches.

According to Dror, 19th-century understanding of ‘emotion’ engaged in a delicate dance with notions of ‘rationality’—while emotion was culturally considered to be outside the scope of rationality, it was brought into the fold of rational cognition in order to be studied in laboratory science. A similar complex dance is currently taking place in HCI. While the ‘turn to emotion’ is often presented as a move away from earlier conceptualizations of rationality and cognition in HCI, we have argued elsewhere that, in fact, it reproduces the very conceptual foundations that it aims to radicalize (Boehner et al., 2005). In what follows, we summarize this argument and how it leads to an alternative model of emotion that may better draw from recent re-understandings of rational cognition in HCI.

2.2. *Affect as information*

Just as existing experimental and methodological frames provided a context by which emotion could be understood

and managed in physiology, so too have these frames guided HCI. Where HCI research has turned to emotion as a relevant concern, it has done so largely within the conceptual framework of rational cognitivism. This is not without some irony. After all, interest in affective computing has been strongly connected with contemporary re-conceptualization of the foundations of interaction (Suchman, 1987) that look beyond traditional cognitive models and in particular turn away from instrumental, procedural, and workplace-focused applications of technology. However, as we shall demonstrate, the very models of cognition as discrete, abstract, and formalizable that are being disbanded for rational thought are often at the heart of how affect is being modeled for computing design. In other words, rather than affect further developing our re-understanding of cognition, affective computing is often following older models that prioritize individual, internal cognition.

Challenges to traditional rational cognition models that neglect or demote emotion and felt experiences have emerged from cultural, sociological, and anthropological critiques as will be discussed in subsequent sections of this paper. In addition, many within the broad fields of cognitive science and psychology, from which this rational model of cognition arose, are also complexifying the construct of emotion beyond traditional dichotomies such as affect vs. cognition, nature vs. nurture, universal vs. local, and conscious vs. unconscious (Lane and Nadel, 2000; Elfenbein and Ambady, 2003; Labouvie-Vief, 2003; Barrett and Wagner, 2006). This work begins to contradict a view of cognition as controlled, precise, and objective and emotion as wild, vague, and subjective. The work of Damasio (1995) in neuroscience, for example, has convincingly demonstrated the interdependence between emotions and activities previously considered to require only rational thought, such as problem-solving and decision-making. The work of Scherer (2005) also argues for models of affect that include subjective feelings in addition to physiological changes and cognitive appraisals. Furthermore, AI researchers, from a field with roots in traditional cognitive science approaches, have a long history of studying ‘emotion’ as a critical aspect of ‘intelligence’ (e.g. Sloman and Croucher, 1981; Elliott, 1992; Bates, 1994; Pfeifer, 1994; Cañamero, 1997), in other words suggesting a more complex integrated model of both intelligence and emotion.

The rise of affective computing, as popularized for HCI by Rosalind Picard and fellow researchers at MIT, begins with this expanded and richer conceptualization of emotions. For instance, Picard (1997) draws on the Damasio studies partly to demonstrate the difficulty in isolating emotions but also partly in her legitimization of emotion as a construct approachable by the sciences. If scientists agree that decision-making and problem solving are worthy pursuits, then the recognition that emotion plays a pivotal role in these processes underscores its importance for understanding both human and computer

intelligence. Similarly, Norman (2004) has become a prominent and influential advocate of emotions as a key component of people’s experience with each other, with the world, and hence with the physical objects around us. Norman, whose studies of design have been hugely influential, has extended his approach to incorporate emotion as a central component, noting that the experience of “everyday things” is conditioned not simply by practical or “logical” concerns but also by aesthetic and emotional ones.

However, although a consideration of emotion as an aspect of interactional experience and competence is presented in these approaches as a move away from traditional cognitive considerations, there is still considerable similarity in how cognition and emotion are treated conceptually in HCI. It is instructive, for example, to compare how cognition and emotion are diagrammed. Fig. 1a is taken from Card et al.’s (1983) classic HCI text, *The Psychology of Human–Computer Interaction*; it shows a schematic overview of their “Model Human Processor,” a quintessential expression of the computational basis of cognition. Fig. 1b is taken from Norman’s (2004) *Emotional Design* and shows an overview of Norman’s three-level model of emotion. What is interesting to note about these two images is where cognition and emotion are located. In both cases, they are contained within the boundaries of the body—caught between eye and hand. Like cognition, emotion is theorized as an internal, thoroughly individual phenomenon.

Both cognition and emotion are construed here as inherently private and information-based: biopsychological events that occur entirely within the body. Like cognition, emotion can be modeled as a form of information processing, and another set of inputs to cognitive processing. This information account of emotion talks about it as a form of internal signaling, providing a context for cognitive action. Emotion as an internal, discrete, and stable phenomenon means it can be measured in an objective, abstract, and standard way. So, rather than radicalizing the informational account of human action implicit in the traditional cognitive model, affective computing often reproduces the very terms of that account.

While accounts of affect are crafted to go beyond cognitive conceptions of mental life, affect is nonetheless understood as the same sort of conceptual object as cognition, and something that might be modeled and made experimentally available in the same sorts of ways.

Like Dror’s 19th century scientists, contemporary affective computing re-inscribes emotion as a dual of rational cognition. As a consequence, sociological critiques of traditional models of cognition that have been influential in HCI take purchase here, as well. In the next section, we draw from anthropology and sociology to present an alternative understanding of the nature of affect. We will move from an idea of emotion as biological, objectively measurable, private and information-based to one of emotion as partially mediated through physiological signals but also substantially constructed through social interaction and cultural interpretation.

2.3. Social and cultural affect

Traditional readings of cognition and rationality have been subject to a continued critique that cognition is relevant and meaningful as a category only in how it is demonstrated and used in the course of everyday social interaction. Scholars such as Schutz (1943) and Garfinkel (1967) draw on a range of empirical material to show that rationality is a witnessable feature of social settings rather than a pure, logical form; the mutual recognition and demonstration of rational behavior is a property of social interaction. In order to understand rationality, then, we must look at the way in which it emerges and is put to work in everyday settings. This is, essentially, an argument about the conceptual categories of cognition and rationality: that they are linguistic terms whose meaning emerges from socially shared practice (Wittgenstein, 1953), so that when we describe the properties of the brain in terms of rationality and cognition, we are in fact re-inscribing features of our social life into our model of mental operation, rather than uncovering features that exist within the phenomena themselves. The idea of rationality—and our interpretation of everyday events as being rational—has a social origin. Drawing on similar foundations, Katz

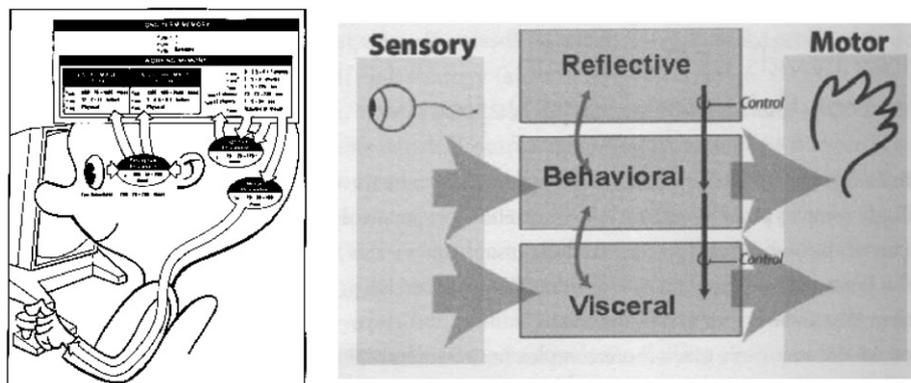


Fig. 1. The model human processor (left, Card et al., 1983) and Norman’s (2004) three-level model of emotion (right).

(1999) has made an analogous argument for emotion, looking at the ways in which emotion is produced in the course of interactions. Where the above informational accounts above, seek an empirical foundation for understanding emotion in measures of biophysiological events, Katz's approach seeks a quite different empirical foundation. Looking at the production of emotion in interaction as the basis for analysis, Katz's empirical basis is, critically, the same as that which underlies our own experiences of emotion as a witnessable feature of daily life.

The importance of such accounts is the fact that they focus on emotion as a social and cultural product. The argument for emotions as cultural has two forms, a weak form and a strong form. The weak form of the argument is simply that emotions have different meanings and valences in different cultural groups. This rests upon a taxonomic conception of "culture" which is itself questionable and which takes an informational account of emotion as a given. Those who subscribe to a taxonomic view of culture believe that understanding a different culture is simply a translation process: with a checklist of properties any culture may be captured, identified and understood. Yet, a dynamic view of culture recognizes it as a living process that escapes easy categorization. Likewise then, the view of emotion as cultural is not about creating a taxonomy of different cultures' emotions.

The strong form of the argument is quite different, and suggests that emotion is not simply interpreted culturally but produced culturally. As an example, consider Lutz's (1988) study of emotion as an aspect of everyday life on the south Pacific atoll Ifaluk. Lutz differentiates here between biological and physiological aspects of feeling, and emotion, which is the culturally grounded set of meanings that both inspires those feelings and gives them lived meaning. Emotion, she argues, is part of cultural and social life. To experience a feeling as, say, anger, love, happiness, lust, or frustration, one must be grounded in a cultural context that makes anger, love, happiness, lust, or frustration meaningful (and in turn determines a response to that emotion—whether it is something to be proud of, ashamed of, etc.).

One example is the Ifalukian concept of *song*, as detailed by Lutz in her ethnographic investigations. *Song* is, loosely, anger. In a Western context, anger is a negative emotion, one that is largely antisocial. *Song*, though, is used rather differently. Lutz translates *song* as "justifiable anger," and notes that its use is, in fact, pro-social; rather than tearing at the social fabric, the cultural use of *song* is cohesive. The justifiable anger of *song* is provoked by a failure to uphold social norms and responsibilities. Taking more than one's fair share at a communal meal, shirking responsibilities in group work, failing to pay appropriate respects to elders or others with whom one stands in a subordinate social relationship, acting inappropriately in social settings, breaking the dignified silence of daily life: these are all actions that might provoke *song*, justifiable anger, in others. Given the strong social shaping to the conditions

under which *song* might be provoked, there is similarly a strong social pressure not to provoke *song* in others. Children are frequently warned against or chided for inappropriate behavior (e.g. boisterous play) by being told that it might make others *song*; similarly, the reluctance to provoke *song* in others is often cited as a justification for particular acts. *Song* is something to be guarded against, and the way in which it is guarded against is by acting in accordance with appropriate cultural conventions.

Song, then, is culturally grounded in two ways. Firstly, the experience of *song* stems from a cultural embedding; it is a response to culturally meaningful events, a personal experience of the violation of norms and expectations which can be understood only with respect to the patterns of cultural interpretation that give social actions meaning. To experience *song*, then, is to be grounded in the cultural patterns that make *song* an appropriate response to have. Secondly, it plays a role in supporting and reinforcing those cultural experiences; the concept of *song* is used to mark behaviors as appropriate or not, as acceptable or not, and so to impose some normative structure on everyday life. Most interestingly, then, *song* has a quite different connotation than anger does in our own culture, due to its pro-social nature; *song* is used to reinforce social structures, patterns, and expectations.

In its strong form, the cultural interpretation of emotion says more than simply that emotions play different roles for different societies, but rather points to emotion as a product and site of cultural meaning-making. In Lutz's analysis, the experience of song is a product of cultural embedding; and similarly, so is Katz's (1999) example of "being pissed off in traffic in LA." Katz's analysis does not simply show that being pissed off in traffic is a uniquely LA experience; rather, he illustrates the ways in which existing patterns of cultural interpretation—issues of age, race, and gender, for example—are enacted in and through the particular moments of emotional interpretations of interaction in traffic. So, Western conceptions of anger, while clearly not pro-social, are still equally culturally situated; they also require an appeal to cultural understandings of the settings within which anger is a culturally appropriate response. That is, the identification of a particular setting (or its associated endocrine reactions) as related to anger (rather than frustration or angst or hatred or disappointment) is every bit as culturally determined as *song*. Neither *song* nor anger is primary, natural, or inherent; they are both cultural products.

This theme runs through a range of studies of the production and management of emotions in different settings, including Geertz's (1957) study of the cultural narratives at work in emotional displays at Javanese funerals; Rosaldo's (1983) discussion of the use of 'shame' as an aspect of the negotiation of individual autonomy, kinship relations and social responsibility amongst the Ilongot of the Northern Philippines; Myers's (1979) study of emotions, selfhood, and morality among the Pintupi; Cahill's (1999) study of emotional work among the

mortuary students and Mann's investigation of the emotional workload of American cocktail waitresses (Spradley and Mann, 1975).

2.4. Interactional and interpretive affect

In summary, then, our argument is two-fold. First, informed by Dror's historical analysis of emotion in physiology, we can see the ways in which the attempt to "tame" emotion as an object of scientific inquiry in HCI has encapsulated it within an existing framework. Like cognition, emotion has been treated as objective, internal, private, and mechanistic. Second, we have suggested that such a reading—what we call the informational reading—systematically ignores a second set of concerns which focus on emotion as it is interactionally and culturally constituted. Again, this is not simply an argument for "cultural difference." Rather, the cultural reading of emotion incorporates two elements. The first is that what constitutes an emotion *at all*—why something is experienced and classified as an emotion rather than as a stomachache, for example—is a cultural question. It is cultural contexts that do or do not allow for such distinctions. The second is that emotion does cultural work. This can be seen, for example, in the ways in which Western conceptions of emotion are embedded within a series of rhetorical oppositions (emotion as hot/cognition as cold, emotions as body/cognition as head, etc), perhaps most significantly the gender association of dispassionate rationality as male and irrational and uncontrollable emotion as female.

We believe that this systematically narrow reading of emotion is problematic for the ways in which we understand emotion in interaction. What we take from the investigations above and related ones that demonstrate historical differences in emotions (e.g. Stearns and Stearns, 1985; Stearns and Lewis, 1998; Gouk and Hills, 2005) is the fundamental principle that an emotion cannot be seen purely as an internal, individual, and private phenomenon. Not only is the experience of emotion mediated by cultural and social situations, it is also used to enact and sustain those settings. As summarized by Schieffelin (1983, p. 181), "the experience, justification, and meaning of affect are not separable from either the role affect plays in the expressive order of interaction, or from the implications of the cultural scenarios in which it participates." Katz (1999) demonstrates the critical role that emotion plays as an aspect of interaction; studies such as those of Lutz suggest, even more strongly, that emotion is an interpretive and interactional product. This is not to dismiss the biophysiological aspects of emotional behavior; rather, we want to understand how those biophysiological events manifest themselves as things that are recognizable, witnessable, collectively and culturally "emotional," and what work is done in the process.

The binding of the social and the cultural, however, does not negate the agency and subjective feelings of the individual. As productive phenomena, culture and social

contexts are also realized, reconstituted, experienced, and over time re-imagined, through the interaction of individuals. With an interactional approach to culture, and subsequently to the experience of emotions, we can reframe the dichotomy of the individual and the social/cultural into a mutually constitutive relationship. In this relationship, emotions are constructed and experienced as individuals act in and through their culture and social interactions. So what we feel is not simply a pre-existing fact, but something that develops over the course of conversations and interactions with one another. We work out what we feel through expressing it and through seeing how others react. We negotiate our feelings with ourselves and with others, over time crystallizing meanings for us of what may initially be vague, confusing, and ambiguous sensations.

From the interactional perspective, affect is not a representational state to be transferred from one place to another, but rather is an aspect of collectively enacted social settings. Emotion is a witnessable property of social action, a way in which actions are rendered interpretable and meaningful. The question of the dynamic, situated interpretation (and attribution) of emotional behavior is critical here. Negotiation, interpretation and inference are inextricably intertwined. Picard (1997) notes that we can never know exactly how someone is feeling, but must always draw inferences about emotional states. Drawing on phenomenological sociology and on McCarthy and Wright's (2004) theories of emotion as part of socially grounded sense-making, we would take this one stage further and suggest that emotion is an intersubjective phenomenon, arising in encounters between individuals or between people and society, an aspect of the socially-organized lifeworld we both inhabit and reproduce. Just as verbal interaction is more than the transmission of information through a conduit, but is rather a form of social action (Clark, 1996), so too is affect a form of social action, both in the ways in which it achieves social ends collectively, and in the ways in which collective meaning shapes individual experience.

3. Designing and evaluating for affect

We welcome the turn to affect in HCI, but have found difficulty in applying standard approaches to affect in ways that address this situated, social notion of emotion. We have found that adopting an information model for emotions rationalizes them, making emotion available to existing scientific models. But rendering emotional experience in informational terms, though methodologically convenient, tends to neglect or obscure a range of considerations that are critical to understanding not just what emotion is but also what it does. We fear that viewing emotion as objectively measurable, biologically constituted facts not only limits but also distorts our understanding of the nature of people's everyday emotional experience. By turning to social, cultural, and interactional accounts of emotion, we hope to enrich our understanding of the

nature of emotional experience. In this section, we highlight the difference the interactional account makes in design and evaluation by providing three pairs of contrasting examples that illustrate how different conceptualizations of emotion lead to different design and evaluation consequences. These examples demonstrate how apparently similar starting points for design can lead to quite different goals, strategies, and implications. In the subsequent section, we will unpack these differences in more detail specifically with regards to evaluation.

3.1. Example 1: *FAIM vs. EmoteMail*

Our first set of contrasting examples concerns adding emotional expressivity to text-based messages such as email or instant messaging, i.e. ways to improve the ‘affective bandwidth’ of text communication. The Facial Affect in Instant Messaging (FAIM) system (El Kaliouby and Robinson, 2005) draws on the informational model to extract emotional information and transmit it to the user’s conversational partner. It uses a video camera to capture a person’s facial expression as he or she is composing an instant message. This image is then analyzed according to indices such as the shape of one’s mouth and head gestures to produce an assessment of one of seven affective states (happy, surprised, agreeing, disagreeing, confused, indecisive, and neutral). This state is subsequently displayed by a cartoon animation accompanying the instant text message. The receiver of the instant message now has an additional cue, namely an animated facial expression, at his or her disposal for understanding the emotional tone of the exchange.

EmoteMail (Angesleva et al., 2004; see also Reynolds and Picard, 2005), on the other hand, is closer to an interactional model; rather than sending symbolic assessments of emotion, EmoteMail sends indicators that can be read emotionally in open-ended ways. EmoteMail takes a series of photos of someone’s face as he or she types an email message. The photo is presented as a black and white thresholded image placed next to the paragraph being typed at roughly the time the photo was taken. In addition, the duration of time spent crafting a particular paragraph is color coded in the text. The receiver of an EmoteMail message therefore also has more cues for interpreting the emotional valence of a message: the abstracted facial expressions and the color-coded paragraphs.

Although these systems appear very similar, the differences between affect as information and affect as interaction play out in a series of different design decisions. In the FAIM example, an emotional model mapping facial features to a limited set of predetermined meanings is built into the system, whereas the snapshot mode of EmoteMail does not include this intermediary classification step and does not restrict the range of expressions. By removing this step, EmoteMail puts the burden of interpreting the facial expressions of the sender on the recipient of the message—something the recipient might be well placed to do if he or

she has previous interactions with the sender to draw from. With EmoteMail the success of the system depends upon it providing new information that lends itself to interpretation and meaning-making by the users. With FAIM, the success of the system is determined by whether or not the meaning of the sender’s facial features is successfully captured and translated into a readable cartoon. The goal of the system is to determine with some degree of accuracy which of seven states the sender is in, capturing this, and then producing simple representations of that state that can be easily decoded by the receiver.

3.2. Example 2: *continuous physiology vs. SEI*

Our second set of contrasting examples addresses how to measure emotion during technology use. These examples focus not on affective systems themselves but on the evaluation of affect in response to systems, in particular electronic gaming systems. Mandryk et al. (2006) do so from an informational perspective through a continuous physiology instrument, which uses a variety of biometric information such as heart rate, galvanic skin response, and electromyography to continuously provide measures of a user’s emotional states while playing a computer game. The instrument works by first establishing a physiological blueprint of different emotional states, such as physiological profiles of boredom, challenge, excitement, frustration, and fun. Once this has been completed, a gamer is wired up with sensor instruments and their physiological responses measured continuously as a game is played. This allows the researchers to identify specific parts in the game where boredom sets in or excitement peaks, for example.

This instrument stands in contrast to the sensual evaluation instrument (SEI) (Isbister et al., 2006) also developed for measuring gamers’ emotional involvement in real-time but drawing on an interactional perspective. The SEI consists of hand-size sculpted objects of various ambiguous shapes, such as pointy edges, bulbous curves, or various combinations of each. The objects are designed to be suggestive of different emotions or feelings, but not directly coded as referring, for example, to ‘boredom’ or ‘excitement.’ It was anticipated that different gamers might use the same object to communicate different feelings. Gamers are asked to pick up and gesture with the objects while playing a video game in order to express how they were feeling. Later, the gamers are asked to explain why they chose certain objects at certain points of game play and reflect on their emotional experiences.

As with our previous set of contrasting examples, the continuous physiology instrument and the SEI start with similar objectives but take very different approaches. For the continuous physiology instrument, the emphasis is on finding and accurately detecting an objectively present, discretely defined emotion. The developers of this method argue, for example, that it fills a gap in existing methods for understanding play environments because it produces quantitative, objective, behavioral measures whereas other

methods such as surveys or questionnaires address attitudes in a subjective and qualitative way. The emphasis in the continuous physiology approach on finding objective indices of emotion illustrates a conceptualization of emotions as a natural fact. For the SEI, the purposeful ambiguity of the sculpted forms draws from and underscores the notion that the emotional information derived from gamers requires additional interpretation. In other words, the shapes in the SEI are not mapped to a limited set of emotional states, but are instead prompts for the gamer and the evaluator to later discuss why certain shapes were chosen as expressive in the moment and what array of feelings were being expressed.

3.3. Example 3: Miro vs. Affector

The last set of examples concerns designing ambient systems for awareness of emotions in an office space, drawing from projects in which two of the authors of this paper were involved. The first system, called Miro (Boehner et al., 2003), drew from an informational model in acquiring and transmitting emotional states. It projected onto the wall in a shared space of an office building a representation of the overall emotional climate based on occupants' real-time responses to an online emotional survey. The display animated an abstract painting, where changes to the background color, the number and cluster of dots, and the speed of movement in the animation reflected aggregates and averages of the survey responses (e.g. feeling happy or sad, tired or wired, etc.). In other words, people's self-reports of feelings, loosely defined, were mapped to changes in the abstract display. In determining what changes should be made in the display, the designers thought about what would make the most intuitive sense to people—for example changing the background color to the yellow end of the spectrum to signify happy and toward the blue end to signify sad (see Fig. 2). As people moved through the common space of the building, they would often stop in front of the display and attempt to read the emotional climate as it had been represented.

The second example, called Affector (Sengers, et al., 2007), draws on an interactional model, shifting from

extraction and transmission of emotional information to providing resources for situated emotional meaning-making. Affector is a video window between the adjoining private offices of two colleagues. The system collects real-time video footage in each office and distorts it based on user-defined rules that map office conditions, such as amount of light, skin color, or movement detected to a desired distortion, such as color inversion, reduction, pixelization, or creating ghost trails of movement (see Fig. 3). The window display in Office A therefore depicts a distorted view of the information collected in Office B and vice versa. Unlike the Miro example, the Affector system explicitly does not include an internal emotional model where the amount of light level in the office is equated with a particular mood and then mapped to a particular video effect. Inspired by Brooks's (1991) argument that systems can appear intelligent and exhibit complex behavior without complex representation and manipulation of abstract information, Affector's behavior may be read as emotionally expressive without it representing emotions internally. In other words, instead of the system interpreting the emotional meaning of the input, the users interpret the emotional meaning of the output, and tune its output to support their readings.

The Miro and Affector system differ in a number of surface ways such as the number of people designed for, the animating display, and the shape of influence over the system (e.g. responding to a survey vs. editing configurable rulesets), but the most substantial difference is in how emotions were conceptualized and experienced. Miro supported an information approach to affect as the system attempted to classify emotional states by the survey, map these states to changes in the display, and then presented the resulting information for the spectators of the display to decode. In order to evaluate this system, we focused on whether the mappings chosen made intuitive sense to people—was it easily readable? Did it seem to capture the 'right' emotional climate? Interestingly, we found the answer to these questions was 'no,' and that Miro failed as emotional information transmission. However, we also noticed that when people encountered the display, they often did so in pairs or groups. In discussing the possible meaning of Miro, they often decoded the display



Fig. 2. Miro display (left) and puzzled users (right).



Fig. 3. Example distortions produced by Affector.

incorrectly, but did so in a way that provided each other with additional background information that Miro could not pick up. For instance, one person said to another: “I think the display is showing a high degree of stress in the office because of all the NSF grant deadlines.” Although according to the underlying information model, this interpretation was wrong, from an interactional point of view the system provided prompts that led these two conversants to move toward a shared understanding of the emotional climate.

This use of a display not as a representation but as prompts to reflection on the affective climate and the dependence on information ‘outside’ the system are both at play in Affector. Affector from the very beginning was intended to be experienced in the conversational way that Miro was appropriated. The designers did not want the system to determine what the affective climate of the neighboring office is but to stimulate reflection on the emotional presence of the other. Much of this reflection and resulting interpretation depends on what users bring to the system, as friends who have existing knowledge of each other and as neighbors who have additional interactions and context information away from the system display which inform their interpretations. In other words, the designers did not imagine that emotions could be captured within the window of the Affector display but rather that the display would become one resource among many in constructing a sense of the other. The evaluation of Affector involved not assessing whether the system captured the emotional climate correctly or whether people were able to read the correct emotional climate displayed.

Instead, the evaluation focused on how the system was engaged for the co-interpretation of affect and how the system attributes encouraged or impaired reflection on the construction of affect.

3.4. Understanding affect as interaction

All three contrasting examples above illustrate the following differences in design and evaluation when approaching affect as interaction vs. information:

3.4.1. The interactional approach recognizes affect as a social and cultural product

With FAIM, the continuous physiology instrument, and Miro, ‘emotion’ is conceptualized as a set of a priori categories that were independent of the concrete context of the system—whereas with EmoteMail, SEI, and Affector, the surrounding context plays a critical role in interpreting affect. For example, with Affector, users undergo a period where they co-construct the affective implications of the system grounded in their existing relationship. Indeed, Affector only works in the context of an ongoing relationship outside of the system that provides the grounds for meaning-making with the system. This shift is present in evaluation as well. For Miro, the designers tried to stand back to observe the system in action, then went away to assess the results from the interviews, conversations, and behaviors. In contrast, for Affector and for the SEI, the evaluators work to present data that is collected as a stimulus for the users, not as a result for the evaluator to analyze.

3.4.2. *The interactional approach relies on and supports interpretive flexibility*

With FAIM, the continuous physiology instrument, and Miro, the ‘meaning’ of the system is intended to be mapped by the designer and ultimately supplied by the system. For example, with the continuous physiology instrument, the evaluators relied on the objectivity of the physiological imprint to help them determine what emotion is experienced, whereas with the SEI, the objects selected by the gamers are more open to interpretation. By leaving the definition of emotion and its interpretation to users, EmoteMail, SEI and Affecter allow emotional meanings to emerge in a situated way over the course of interaction.

3.4.3. *The interactional approach avoids trying to formalize the unformalizable*

In informational approaches, emotions are generally characterized as discrete units or as points in a multi-dimensional space—such as the defined set of five or six emotional states used in FAIM, the continuous physiology instrument and Miro. Even when they are based on psychological models of emotion, such as with the continuous physiology instrument, it is not clear that emotions as experienced by users in complex social and cultural ways map neatly onto these underlying emotional structures. Sometimes emotions cannot be articulated by users in straightforward ways, yet informational approaches can unintentionally attempt to force users into a straightjacket of formalized expression. The interactional approach does not require emotion to be formalized by the system or in evaluation; instead, the users can supply the emotional meaning in the system.

3.4.4. *The interactional approach supports an expanded range of communication acts*

Informational approaches focus on communicating affect through a well-defined set of signs with clear meanings assigned by the designer of the system—for example, in the case of Miro, “fast-moving dot” should mean “high energy office” whereas “slow-moving dot” should mean “lethargic office.” In Affecter, meaning is communicated through a combination of video feed and a distortion language that can be overlaid on that feed. Rather than designer-defined signs, Affecter supports user-defined signs as well as indexes and icons, which can give a more open-ended sense of the complexity of emotion.

Similarly, in systems like EmoteMail or the SEI, emotion can be communicated in a richer way than clearly defined signs allow. Instead of users thinking, ‘I feel excited. The pointy object represents excitement in this system. Therefore, according to the system language, I feel pointy,’ an alternative is to allow users to express themselves directly or to use the expressive capabilities of the system to demonstrate, rather than represent, their feelings: ‘I’m feeling pointy. Pointy demonstrates how I’m feeling right now.’ In this example, the exact interpretation of what ‘pointy’ means is left open to the people involved and

depends on the detailed situation of their discussion. A similar approach is followed in the design of eMoto, a system designed for expression of affect in mobile phone text messages (Sundström et al., 2005). eMoto allows users to alter the background color and pattern of their message through affective gestures with the stylus. Rather than requiring a verbal articulation and translation of how someone is feeling, the system allows users to shake it out, to demonstrate with varying pressure, movement patterns, and pace something that reflects how they feel.

3.4.5. *The interactional approach focuses on people using systems to experience and understand emotions*

Given the complex, ambiguous, and situated nature of emotions, it seems unlikely that emotions will ever be fully understood by computer systems. The interpretational approach side steps this problem, since the focus is instead on using systems to stimulate reflection on and awareness of affect. For example, in Höök et al.’s articulation of the affective loop (e.g. Fagerberg et al., 2003), affective input to systems is managed, not by extracting emotional information from users, but by having users directly express emotions to systems. While users may express emotions they do not feel, in an affective loop the expressive gestures and the system’s reactions are set up to reinforce whatever emotion the user expresses. That is to say, a user may express an emotion they do not feel, but they will likely come to feel the emotion they are expressing through the course of the interaction. Such a system does not primarily transmit emotion as information but instead supports its experience and construction. The important thing from the interactional perspective is not making systems more aware of emotions but making people more aware of emotions through system use and design.

4. Making by measuring

The differences described above between thinking of affect primarily as information and considering affect as centrally enacted in interaction substantially change the playing field for system design. Instead of designing systems to decode and encode affective signals, systems are set up to engage users in interactions that help them to understand, reflect on, and experience their emotions in new ways. In the process, as we have shown, evaluation changes as well. One obvious change is in the metrics of evaluation, since the goals of these new systems are substantially different. But deeper issues are involved, ones that suggest substantial changes in the constitution and interpretation of evaluation methods.

In this section, we will argue that the dominant strategies used for evaluating affective computing are grounded in the notion of affect as information. In the process, they re-inscribe affect as necessarily informational. In other words, currently common evaluation methods seek a particular kind of emotion, and, in finding it, reinforce the conception of emotion that they began with. If we

instead take affect as interaction, standard evaluation methods need to be substantially retooled in order to uncover the sorts of emotional experiences this new model takes as central. We will demonstrate here three major differences that underlie these different perspectives on evaluation: whether they take emotion to be delineable or inherently ambiguous, whether they seek objective measures or subjective experiences, and whether emotion is adjudicated by the evaluator or jointly constructed by evaluators and subjects. In the following section, we will describe how these differences lead not only to different methods for evaluation, but also to very different conceptions of what evaluation is in the first place.

4.1. *Delineable vs. ambiguous*

The first set of differences between these approaches to evaluation is in whether emotions are considered to be well defined or ambiguous, i.e. fundamentally beyond the reach of categorization. The informational model is based on a notion of emotions as delineable, and evaluation methods used to identify such emotions both draw on and reinforce the assumption that emotions are clearly defined. For example, a conception of emotion as a discrete phenomenon leads to the use of evaluation methods such as the label approach with a limited set of categories, either pictorial or words, for users to identify their emotions. Likewise, with a dimensional view of emotions, instead of asking users to choose how they are feeling from a set list of emotions, users are asked to place their feeling in a 2- or 3-dimensional grid according to, for example, how aroused or how positive or negative they feel. While there are differences between these methods—dimensional methods allow for a greater range of expressions, and the use of dimensions instead of labels suggests that emotions slide into and overlap with each other—both methods function to turn experienced and perhaps ambiguous emotions into quantified, delineable categories. Both these methods are applied in order to provide metrics that make emotion clearly classifiable and comparable within and across individuals (Scherer, 2005).

Similar issues arise with how subjects are chosen and managed in order to generate recognizably delineable, standardized emotions. For example, tests of a computer's ability to detect emotions often take place in the laboratory with actors demonstrating scripted emotions. These tests are based on the assumption that the kind of exaggerated, clearly delineated emotions that actors express correspond directly to emotional experience “in the wild.” Dror (1999b) notes a similar management of users in his historical discussion of the numerization of emotions. In order to measure emotion, Dror noted that scientists chose subjects who produced clearly recognizable emotions on demand, either physiologically or through facial gestures. In other words, subjects who could communicate emotions demonstrably, as an actor would, were preferred over those who did not, resulting in experiments with clearly delineated emotions.

As a more recent example of this historical disjuncture between emotions found in the laboratory and those “in the wild,” Batliner et al. (2003) review the performance of automatic call-in systems that use experimental pre-tests of the software's ability to detect frustration in a caller's voice. In the laboratory, the systems performed with a high success rate in response to the actor's exaggerated portrayal of emotion on cue. But, in the field, the systems performed poorly. The difference could be attributed to the assumed transferability designed for in the experimental setting. The actor has a script and a set emotion to convey to the system for detection—more than likely the stereotypical cues the actor would use to convey frustration are the same cues used to design the system itself, therefore ensuring a high degree of success. But emotion on the ground is more subtle, complex, and ambiguously presented.

What evaluation methods grounded in the assumption of delineability foreclose is the potential ambiguity of emotional experience: the possibilities, for example, that subjects may be feeling multiple emotions at once, or that the nature of the emotions felt may simply be unclear. Methods that ask subjects to classify an emotion exactly necessarily turn these experiences of ambiguity into apparently clear categories. Approaches that disqualify particular subjects or particular settings because the emotions felt are not clear or extreme enough close their eyes to the complexity and subtlety of the everyday emotional practices around computation.

Methods that derive from the model of emotion as interaction, on the other hand, seek to draw out and understand the ambiguities inherent in emotional experience. Such methods prioritize the complex expression of the experience of emotion over its categorization into delineable categories. Some of these methods draw on similar techniques to those used in the informational model; one could, for example, present users with a set of emotional categories, but use these as a foil to discuss their experiences (as in “I feel a bit unhappy, but it's not a regular kind of unhappiness...”) rather than as a taxonomy to classify them. Others steer clear of categorization altogether, drawing on interviewing, observation, and self-report to generate open-ended reports of emotional experience in action. The evaluation of eMoto, for example, as described by Sundström et al. this issue, drew on open-ended methods such as interviewing and recruiting friends who were intimately acquainted with each others' emotional experiences to make ‘experience clips,’ or brief movies of personally meaningful events involving eMoto. These methods also suggest the central value of subjective measures of emotion, a point of difference between informational and interactional models we explore in more detail next.

4.2. *Objective vs. subjective*

The second set of issues we will explore focus on which kinds of measures are sought: objective measures which can report user emotional experience without the user's

intervention, or subjective measures based in part on self-report. While many affective evaluations effectively use mixed measures, i.e. a combination of objective and subjective measures, there are differences based on which of those measures is seen as most reliable or desirable.

One of the areas in which this tension can be seen most clearly is in the growing popularity of physiological metrics for measuring emotion. Their current uptake does not derive from the fact that they are new—many of these techniques have been around since the early 1900s (Dror, 2001)—but from their appropriateness for a perspective on emotion as a delineable numeric value of an internal phenomenon, which therefore may be correlated to biological signals. The physiological measures are valued for being unobtrusive and therefore able to monitor continuous indices of emotion, rather than asking the user to stop their experience, reflect on their emotion, and then disclose it. Since emotion is seen as a biological fact that pre-exists interaction or interpretation, these physiological measures are said to represent empirical evidence of how a user truly feels, before he or she has the opportunity to mask this feeling or to reason it away.

Yet the growing popularity of such metrics is based not only on the attractiveness of physiological signals as empirical measures of a natural phenomenon, but also on a suspicion of what are considered more subjective measures. In other words, it is not just that researchers are selecting objective physiological measures; they are also actively choosing *not* to select alternative measures such as self-report. Although many researchers advocate a mixed-methods approach, there is often a clear preference for the objectivity of physiological measures, and therefore presumed scientific superiority, over the subjectivity of self-report (see for example Mandryk et al., 2006).

Despite this preference, however, systems of measurement that prioritize physiological signals as direct evidence of emotion still depend crucially on subjective measures, since their validity often depends upon gauging the meaning of those metrics against subjective reports. For instance, in determining that heart rate is a viable metric for assessing the positive or negative valence of an experience, the protocol involves first discovering a pattern (e.g. heart rate predictably going up or down in response to subjectively defined stimuli such as pictures from the International Affective Picture System (Lang et al., 2005)). Once a pattern is established and correlated with the emotionally valenced stimuli, this correlation must be validated by asking subjects if they would classify their emotion as positive or negative. In order for the physiological measurement to be considered valid, it must first be shown to match the self-assessment of subjects.

Once this correlation is established, however, subjective measures come to be replaced with physiological measures. The physiological measures then stand for the subjective accounts as a more reliable and continuous source without needing the conscious involvement of the subject. The

movement to validate objective measures against subjective ones, then, is followed by a second movement that erases its tracks, where the subjective accounts are now positioned as inferior. Although the subjective accounts stand at first as the standard or benchmark from which to develop the physiological measures, the second move identifies the limitations of self-report as weaknesses. For instance, self report of emotions is held as suspicious because subjects may not remember how they actually felt, they may try to misrepresent how they felt due to self-presentation or trying to please the researcher, they may not be able to articulate how they feel, or if they can articulate their feelings they often describe them in a non-extreme manner, therefore making statistical variance difficult (Picard and Daily, 2005).

These are all presented as limitations of the evaluation method of self-report when the affect one is aiming to measure is considered discrete, pre-defined, transferable, and self-contained. When emotion is found to be different, this is considered not evidence for another kind of affect but evidence that the methods are flawed. In this way, the prioritization of physiological evaluation methods not only supports the discovery of apparently objective measures of emotion, but also perpetuates a view of emotions as objective phenomena in the first place.

While evaluations that draw from affect-as-information see emotion as an objective, natural fact which is distorted in reflection and expression, evaluations that draw from the affect-as-interaction model see emotion as constructed through interaction and expression, and distorted if presented *without* this intermediary process of meaning-making. From this perspective, physiological signals alone are not enough to understand the subjective experience of emotion, since emotions as experienced constitute more than physiological responses, including layers of meaning-making and reflection. Indeed, since physiological signals tied to emotion may be overloaded with ones tied to other aspects of behavior, users may or may not even experience the same signals as emotional in different contexts.

More broadly, in taking an interactionist approach to emotion, it becomes clear that overdependence on evaluating what can be measured objectively misses and discounts the subjective experience of emotion. It is this very subjectivity that defines affect in an interactional approach. In this model, the shifting constructions of affect in self-report are not limitations at all, but evidence of how emotion is socially, culturally, and relationally developed. Therefore, many of the interactionist evaluation methods specifically attempt to draw out subjective experiences of emotions. This is certainly not to the exclusion of external metrics such as physiological measures; as we will describe below, such measures are useful as prompts for eliciting subjective interpretations. We simply recognize that using objective measures to the exclusion of subjective accounts in evaluation leads to a too-narrow, often misleading understanding of affect.

4.3. *Adjudicated vs. co-interpreted*

In the previous two sections, we have unpacked two ways in which the evaluation methods for affect as information differ from those for affect as interaction. The difference between these models runs deeper than the methods used, however; the models disagree at a fundamental level about what it means in the first place for an affective computing system to succeed.

The informational model is based on a notion of emotion as transferable, communicable units. Systems drawing on this model are thought of as transferring information in ways analogous to Shannon's (1949) theory of information, i.e. a sender transmitting an emotion to a receiver through a noisy medium. Success of such a system can be measured by whether the sender's original emotion is the same as that understood by the receiver, whether the sender or receiver is another user or a computer. The focus of evaluation, then, is on identifying the correct emotion in the sender and receiver. If the same emotion is sent, processed, and received, the system can be said to be a success.

Evaluation therefore depends crucially on the evaluator's ability to convincingly identify the single, correct emotion of human users, using methods such as those previously described. This, again, can be thought of as a problem of information transmission: the sender (the user) transmits their emotion through the evaluation method (e.g., physiological sensing) to the receiver (the evaluator). A good evaluation method will assure that the evaluator received the correct original emotion, undistorted by noise such as the user's possibly incorrect claims about the nature of his or her experience. In doing so, the evaluator is set as the final authority on the user's true emotion.

The interactional model instead sees emotion as constructed through interaction and expression. The role of affective systems is not to transmit pre-existing emotional units, but to provide a resource for emotional meaning-making. Success of such a system is measured by whether users find the system's responses useful for interpreting, reflecting on, and experiencing their emotions. Evaluation does not aim at finding a user's original, 'true' emotions, but in tracking how emotions are constructed and interpreted over time, and correlating these dynamics with aspects of system design. Just as emotion is constructed in interaction with systems, so it is constructed in interaction with the evaluator. The evaluator does not and cannot adjudicate which is the user's correct or actual emotion; instead, he or she provides stimuli that elicit emotional reactions and reflections. Evaluators and users co-interpret the results, together constructing an account of what happened and why—perhaps even along with other parties as in Gaver's contribution to this issue. These accounts may not converge on one, single, correct interpretation but instead leave open the opportunity to play those interpretations off of each other in developing a rich account of system use (Sengers and Gaver, 2006).

The difference between these models can be demonstrated by returning to the comparison between Mandryk et al.'s (2006) continuous physiological method, in which physiological signals are collected during game play and used by evaluators to identify emotionally salient moments, and Isbister et al.'s (2006) SEI, in which players consciously select and later reflect on shapes that represent their feelings during the course of game play. The difference in these methods is perhaps most striking in the variety of the results they generate. The heart-rate measure produces a graph of beats per minute matched to the various episodes in the game. This evidence can be supplemented with the gamers' accounts of what is going on, but foregrounds the numerical representation of value. With the SEI method, the gamers' selection of different ambiguous forms generates a variety of responses and approaches to the shapes and the instructions for reflection. The stimulus in this case is still an abstract form, but one without a preset meaning. The heart-rate measures result in what Dror (2001) refers to as a tradable and easily comparable representation of emotion, while the SEI results in personal and idiosyncratic representations.

Furthermore, in the heart-rate model, the information is generally not shared with the user, lest it influence the measurements. The evaluator tends to be the only one privy to the physiological measures, whereas with the SEI, the evaluator and the user collaboratively explore what the different shapes might mean. The goal in the former is to autonomously identify the user's correct emotion; the goal in the latter is to co-interpret it. The former follows from and confirms a view of affect as information, where the evaluator is the final judge, whereas the latter follows from and confirms a view of affect as interaction, where evaluators and users collaborate to better understand user emotional experience.

Again, this does not mean that physiological or other behavioral and empirical measures are not useful when taking emotion as interaction. Existing measures developed in the affective computing community are a powerful tool for interactionist approaches as well, where they are used collaboratively to develop an understanding of the users' subjective experience. In our evaluation of *Affector*, for example, we track objectively measurable data such as number of interactions and use it to elicit narrations of subjective experiences. Physiological measures such as heart-rate information could be an interesting feedback mechanism that stimulates users themselves to narrate why their heart rate is increasing or decreasing. This illustrates the fundamental difference in measuring affect as something to be found by the researcher or something that is constructed in practice.

5. Beyond evaluation?

So far, our story has been one of symmetric alternatives. Emotion can be taken as information, or it can be taken as interaction. Both models of information can usefully

inform design and evaluation. Both approaches seek to find a particular kind of emotion, and both succeed in finding it. There is, however, one crucial difference. Evaluating affect-as-information constructs its object of study, i.e. emotion as delineable and transmissible units, but takes that kind of emotion to exist independently of its methods. Of course, affect-as-interaction constructs its object of study as well; but it takes that construction as a necessary part of its practice. In other words, affect-as-interaction highlights the role evaluators play in creating the kind of emotion they are studying, while affect-as-information tends to view evaluators as uncovering an objective truth. What the focus on categorization, purported objectivity, and adjudication in affect-as-information sidelines is the ways in which evaluators and subjects in *both* approaches are inextricably intertwined in constructing the emotions that are being measured. This raises serious questions for all parties for what it means to manipulate, measure, or interpret emotional responses as part of interactive systems practice.

We find it useful to think about this problem in terms of “epistemic regimes”—systems of practice and regulation that produce and reinforce commitments to particular forms of knowledge and knowledge production. The relevance of this concept here is that it can help us to see how problems arise within established disciplines as they attempt to incorporate new phenomena, because these new phenomena need to be transformed so that they can be handled with existing conceptual and methodological tools.

So, for example, the natural science model that underlies conventional HCI design practice—and most particularly its approach to evaluation—is one in which the observable phenomena of interaction are taken as indicative of internal, hidden events that are foundational, causal, efficacious, and can be modeled in mathematical and informational terms. The problem is, then, to separate the “signal” from the “noise.” In Picard and Daily’s (2005) discussions of self-report, for example, the problems of forgetting, of normalization, and of willful misrepresentation are identified as ways in which the “true”—internal—phenomena are obscured.

The historical material presented earlier vividly demonstrates the problem of epistemic regimes, in which physiologists first recognized the relevance of emotion for their work and then incorporated it into existing models of scientific investigation and laboratory practice. What is perhaps most telling here is the way that “unruly” emotion is disciplined as a practical matter and as a conceptual one—emotion becomes measurable because it is the only way to imagine treating it scientifically. We would suggest that current approaches to considering emotion as a component of HCI evaluation exhibit similar features.

Evaluation depends on measurement, on comparison, and, principally, on expectation. One can evaluate only with respect to a goal. When we look at emotion as a product of culturally situated encounters between people and settings of action, measurement and comparison become problematic, but most particularly our notion of

expectation begins to fail us. In all that we have presented, culture is a generative rather than taxonomic concept—something that “gives rise to” rather than something that “categorizes.” Accordingly, the idea that emotion is a cultural product not only weakens (at best) physiological models as analogs of emotional experience, but also suggests that the distance between settings of design and settings of use questions the position of designer expectations within the framework of interactive systems. When we talk of “epistemic regimes,” then, we draw attention to the ways in which particular forms of knowledge regiment, organize, and discipline not only design processes but technology use. In particular, we want to note the ways in which particular approaches to evaluation (as above) are shaped by the kinds of epistemic apparatus available. The central problem here is that evaluation, as an epistemic regime, frames emotion as an analog of performance, efficiency, and technical achievement within interaction. The production of emotion is a goal, and what is evaluated is not emotion in and of itself, but rather the effectiveness with which that goal can be achieved. In support of measurement, emotion must be made measurable and accountable. In this model, observable phenomena are useful only in as much as they provide access to internal mechanism.

So, in drawing on self-report in our own work, we are not simply attempting to patch holes in a methodological framework, but rather drawing on a radically different model. The alternative model that we have suggested—in which emotion is itself a cultural and interactional product, and self-reports can therefore be scrutinized as sites where emotion is produced—is not simply a different base for empiricism, but a quite different perspective on the production and validation of scientific knowledge. A cultural account of emotion places a very different emphasis on the observable features of everyday life. Interaction is the site where emotion is produced, rather than its consequence; attempting to push past interaction in order to uncover the emotion beneath is, as Garfinkel (1967, p.22) memorably suggested, “very much like complaining that if the walls of a building were only gotten out of the way one could see better what was keeping the roof up.”

The question that this raises is quite what “evaluation” is, and what we might like it to do for us. One strategy is to radically revise the notion of evaluation so as to make it supportive of alternative forms of knowledge and knowledge production. From this perspective, we might broaden the definition of evaluation in the hope of making it amenable to different sorts of epistemological commitment. This seems, though, a quixotic task. A more practical alternative is to focus less on what evaluation is and more on what it does. Evaluation, we would argue, is one among a range of strategies of critique. As Höök et al. (2003) suggest, traditional HCI evaluation methods can be brought to bear on interactive systems in which emotion plays a central role, but will inevitably fail to address the

full range of considerations by which a system might be judged not simply effective but interesting or successful. Just as we have suggested that different perspectives on emotion can lead to alternative design strategies, we need also to incorporate different models of critique into the design lifecycle.

6. Conclusion

Emotions occur dynamically at the interface of experience in the world. As Dewey remarked,

Joy, sorrow, hope, fear, anger, curiosity, are treated as if each in itself were a sort of entity that enters full-made upon the scene, an entity that may last a long time or a short time, but whose duration, whose growth and career, is irrelevant to its nature. In fact emotions are qualities, when they are significant of a complex experience that moves and changes ... All emotions are qualifications of a drama and they change as the drama develops. (Dewey, 1934, as cited in McCarthy and Wright, 2004, p. 83)

This perspective of emotion as strongly situated, moving and changing, entering the scene incomplete, is difficult to reconcile with approaches to emotion as one of transmission of information. We have been concerned here with an alternative reading of affect, in line with Dewey's concerns about reading emotion as natural facts and as discrete events. We have examined an alternative account of emotions as interactional products rather than informational objects, and demonstrated how this approach supports an alternative form of affective computing design. This reading draws from and further develops other alternative approaches to affective computing (e.g. Fagerberg et al., 2004; Gaver et al., 2004; McCarthy and Wright, 2004; Sundström et al., 2005).

More broadly, this is part of a larger research program into the relationship between technology and practice, and in particular a move from technologies of representation to technologies of participation. Participation emphasizes the ways in which information systems act as platforms upon which social structure is enacted, rather than as entities employing representations of the world and therefore always at one step removed from it. Emotion is a particularly interesting topic from this perspective, precisely because it is deeply enmeshed in a broad range of cultural meanings and oppositions—mind/body, cold/hot, male/female, serious/frivolous, etc. The complexity and dynamism of emotion that Dewey observes are precisely the properties that make emotion suitable for this treatment. The richness of emotion in interaction mitigates against reductive representation.

Furthermore, as information technologies increasingly inhabit the everyday world, we need to understand them on multiple levels simultaneously—as technological artifacts, social facts, and cultural narratives. As we have

demonstrated here using emotion as a lens, this view is not a technologically limiting one, but rather opens up new avenues for design and evaluation.

While much work in affective computing is grounded in the laboratory behavioral sciences, we have been inspired primarily by anthropological and historical approaches to emotion. This has led us to move from a model of emotions as naturally existing and objectively measurable to one of emotions as interactionally constructed and subjectively experienced. But, as we have discussed in the last section, if affect is not only measured but also made in interaction, the idea that we could or should study emotion only in the manner of a natural science—as a phenomenon that is objectively there, irrespective of the actions of the researcher—becomes problematic.

We have argued that physiological approaches to emotion, in focusing on measuring emotion as objective and well defined, fail to address how emotions are actually experienced. This is not necessarily problematic for the laboratory behavioral sciences, where it may be more useful to develop reliable, reproducible hypotheses around emotion than to fully reflect the messiness and complexity of people's lived experiences of emotion. The story is different for computing design, however. Given the pervasiveness of computing technology in our everyday lives and its concomitant societal impact, it is essential that we address people's actual lived emotional experiences around technology. This requires us to move beyond natural-scientific conceptions of emotion, which tend to sideline people's felt emotions in preference for objective measurements.

Debates around the nature of affect—as subjectively experienced, or as objectively measured—have existed as long as emotion has been studied. In 1962, the pioneer of affect psychology Silvan Tomkins wrote, “[t]he argument is whether a human being is the most real and valuable entity in nature, an end in himself—or whether that which is most real and valuable is some norm, some ideal essence quite independent of the human being. Those who have insisted on the objective, whether in science or in art, have insisted on certainty. Those who have insisted on the subjective have glorified man and stressed the value of both play and risk. It is our conviction that the humanistic ideology, in the long run, will yield both the greater and the surer payoff.” (Tomkins, 1995, pp. 202–203). Tomkins's sentiment suggests a view of certainty as located in the authenticity of the subjective. Although our approach resonates with this view, we argue that the way forward for affective computing and affective evaluation is not a debunking of objective approaches in total but a recognition of the limits and liabilities of both objective and subjective accounts of emotion. We advocate a need for greater appreciation of the interplay between these accounts and their underlying assumptions, for example how an observable phenomenon such as heart rate becomes identified as an objective measure of emotion in the first place and how the value of different accounting

methods are determined. This recognition leads to an appreciation not only of how emotion as a phenomenon is measured but also how measures make the phenomenon of emotion itself.

What we offer, finally, is not a refutation of affect as information by affect as interaction. Instead, we show that a very different conception of emotion from the standard model can be usefully embodied in technology design and evaluation. What we offer, then, is choices, based on the realization that there are multiple ways to understand what emotion is, and multiple ways to approach it usefully in technical practice. In the process, we aim to short-circuit the mechanisms by which the informational model's embodiment in technology and analysis through evaluation has come to warrant its apparent inherent truth.

Indeed, understanding how different perspectives on emotion shape what is found in experimental practice calls into question the standard reasons given for choosing informational methods. A principled decision of which methods to use cannot rest simply on the clarity of the emotions they find, since this clarity rests in the methods chosen, and is not an attribute of the underlying phenomenon they seek to approach. It cannot rest simply on the apparent objectivity of the approach, since this objectivity is grounded in and validated by subjective methods. It cannot rest simply on the ability to give the one correct answer about what a user feels, since the correct answer is not necessarily also true.

Each model provides a different lens on emotion; neither can be said a priori to be more correct. Given the fact that we cannot ground our decision on which gives direct access to emotion, we would suggest that decisions about what choice to take should take into account the political, social, and cultural implications of holding a particular view on emotion. We need to ask ourselves who gets to decide what emotion a person feels, what emotional experiences are designed for and which are left out, and whether the resulting technologies enrich or flatten emotional experience. Correct answers to these questions, in the end, matter more than the correctness by which a computer can judge how a person feels.

Acknowledgments

This work was funded in part by NSF Grants IIS-0133749, IIS-0205724, IIS-0238132, IIS-0326105, and IIS-0534445, and by Intel Corporation. We are grateful for the inspiration and support of many colleagues, including Ken Anderson, Genevieve Bell, Johanna Brewer, Richard Doyle, Maria Fernandez, Brooke Foucault, Bill Gaver, Geri Gay, Kristina Höök, Katherine Isbister, Joseph “Jofish” Kaye, David Klein, Scott Mainwaring, Michael Mateas, Jay Melican, John McCarthy, Yevgeniy “Eugene” Medynskiy, Eunyoung “Elie” Shin, Petra Sundström, Simeon Warner, Elizabeth Wilson, and Peter Wright.

References

- Angesleva, J., Reynolds, C., O'Modhrain, S., 2004. Emotemail. In: Proceedings of the 31st International Conference on Computer Graphics and Interactive Techniques.
- Barrett, L.F., Wagner, T.D., 2006. The structure of emotion: evidence from neuroimaging studies. *Current Directions in Psychological Science* 15 (2), 79–83.
- Bates, J., 1994. The role of emotion in believable agents. *Communications of the ACM* 37 (7), 122–125.
- Batliner, A., Fischer, K., Huber, R., Spilker, J., Noth, E., 2003. How to find trouble in communication. *Speech Communication* 40 (1–2), 117–143.
- Boehner, K., Chen, M., Liu, Z., 2003. The vibe reflector: an emergent impression of collective experience. In: CHI 2003 Workshop on Providing Elegant Peripheral Awareness.
- Boehner, K., DePaula, R., Dourish, P., Sengers, P., 2005. Affect: from information to interaction. In: Proceedings of the Critical Computing, pp. 59–68.
- Boellstorff, T., Lindquist, J., 2004. Bodies of emotion; rethinking culture and emotion through Southeast Asia. *Ethnos* 69 (4), 237–444.
- Brooks, R.A., 1991. Intelligence without representation. *Artificial Intelligence* 47 (1–3), 139–159.
- Cahill, S., 1999. Emotional capital and professional socialization: the case of mortuary science students. *Social Psychology Quarterly* 62, 101–116.
- Cañamero, D., 1997. Modeling motivations and emotions as a basis for intelligent behavior. In: Proceedings of the First International Conference on Autonomous Agents, pp. 148–155.
- Card, S., Moran, T., Newell, A., 1983. *The Psychology of Human-Computer Interaction*. Lawrence Erlbaum, Hillsdale.
- Clark, H., 1996. *Using Language*. Cambridge University Press, Cambridge.
- Damasio, A., 1995. *Descartes' Error: Emotion, Reason and the Human Brain*. Quill, New York.
- Dewey, J., 1934. *Art as Experience*. Perigree, New York.
- Dror, O., 1999a. The affect of experiment: the turn to emotions in Anglo-American physiology, 1900–1940. *Isis* 90 (2), 205–237.
- Dror, O., 1999b. The scientific image of emotion: experience and technologies of inscription. *Configurations* 7 (3), 355–401.
- Dror, O., 2001. Counting the affects: discoursing in numbers. *Social Research* 68 (2), 357–378.
- El Kaliouby, R., Robinson, P., 2005. Generalization of a vision-based computational model of mind-reading. In: Proceedings of the First International Conference on Affective Computing and Intelligent Interaction, Beijing, China, pp. 582–589.
- Elfenbein, H.A., Ambady, N., 2003. Universals and cultural differences in recognizing emotions. *Current Directions in Psychological Science* 12, 159–164.
- Elliott, C., 1992. *The affective reasoner*. Ph.D. Thesis, Report #32, Northwestern University, Institute for the Learning Sciences Technology.
- Fagerberg, P., Ståhl, A., Höök, K., 2003. Designing gestures for affective input. In: Proceedings of the Mobile Ubiquitous and Multimedia (MUM).
- Fagerberg, P., Ståhl, A., Höök, K., 2004. eMoto: emotionally engaging interaction. *Personal and Ubiquitous Computing* 8 (5), 377–381.
- Garfinkel, H., 1967. *Studies in Ethnomethodology*. Polity, Cambridge.
- Gaver, W., Boucher, A., Pennington, S., Walker, B., Bowers, J., Gellerson, H., Schmidt, A., Villars, N., Steed, A., 2004. The Drift Table: Designing for Ludic Engagement. *Extended Abstracts of CHI 2004*, pp. 885–900.
- Geertz, C., 1957. Ritual and social change: a Javanese example. *American Anthropologist* 59 (1), 32–54.
- Gouk, P., Hills, H., 2005. Towards histories of emotions. In: Gouk, P., Hills, H. (Eds.), *Representing Emotions: New Connections in the Histories of Art, Music, and Medicine*. Ashgate, Burlington, pp. 15–34.

- Höök, K., Sengers, P., Andersson, G., 2003. Sense and sensibility: evaluation and interactive art. In: *Proceedings of CHI 2003*, pp. 241–248.
- Isbister, K., Höök, K., Sharp, M., Laakso, J., 2006. The sensual evaluation instrument: developing an affective evaluation tool. In: *Proceedings of CHI 2006*.
- Katz, J., 1999. *How Emotions Work*. University of Chicago Press, Chicago.
- Labouvie-Vief, G., 2003. Dynamic integration: affect, cognition, and the self in adulthood. *Current Directions in Psychological Science* 12 (6), 201–206.
- Lane, R., Nadel, L., 2000. *Cognitive Neuroscience of Emotion*. Oxford University Press, Oxford.
- Lang, P.J., Bradley, M.M., Cuthbert, B.N., 2005. International affective picture system (IAPS): digitized photographs, instruction manual and affective ratings. Technical Report A-6. University of Florida, Gainesville, FL.
- Lutz, C., 1988. *Unnatural Emotions*. University of Chicago Press, Chicago.
- Mandryk, R., Atkins, M.S., Inkpen, K., 2006. A continuous and objective evaluation of emotional experience with interactive play environments. In: *Proceedings of CHI 2006*.
- McCarthy, J., Wright, P., 2004. *Technology as Experience*. MIT Press, Cambridge, MA.
- Myers, F.R., 1979. Emotions and the self: a theory of personhood and political order among Pintupi Aborigines. *Ethos* 7, 343–370.
- Norman, D., 2004. *Emotional Design: Why We Love (or Hate) Everyday Things*. Basic Books, New York.
- Pfeifer, R., 1994. The “Fungus Eater” approach to emotion: a view from artificial intelligence. *Cognitive Studies* 1, 42–57.
- Picard, R., 1997. *Affective Computing*. MIT Press, Cambridge.
- Picard, R., Daily, S.B., 2005. Evaluating affective interactions: alternatives to asking what users feel. In: *CHI 2005 Workshop on Innovative Approaches to Evaluating Affective Systems*.
- Reynolds, C., Picard, R., 2005. Evaluation of affective computing systems from a dimensional metaethical position. In: *Proceedings of the First Augmented Cognition Conference, in conjunction with the 11th International Conference on Human-Computer Interaction*.
- Rosaldo, M., 1983. The shame of headhunters and the autonomy of self. *Ethos* 11 (3), 135–151.
- Rosaldo, M., 1984. Towards an anthropology of self and feeling. In: Sweder, LeVine (Eds.), *Culture Theory: Essays on Mind, Self, and Emotion*. Cambridge University Press, Cambridge.
- Scherer, K., 2005. What are emotions? And how can they be measured? *Social Science Information* 44 (4), 695–729.
- Schieffelin, E., 1983. Anger and shame in the tropical forest: on affect as a cultural system in Papua New Guinea. *Ethos* 11 (3), 181–191.
- Schutz, A., 1943. The problem of rationality in the social world. *Economica* 10, 130–149.
- Sengers, P., Gaver, B., 2006. Staying open to interpretation: engaging multiple meanings in design and evaluation. In: *Proceedings of DIS 2006*.
- Sengers, P., Boehner, K., Mateas, M., Gay, G., 2007. The disenchantment of affect. *Journal of Personal and Ubiquitous Computing*, in press.
- Shannon, C., 1949. *The Mathematical Theory of Communication*. University of Illinois Press, Urbana.
- Sloman, A., Croucher, M., 1981. Why robots will have emotions. In: *Proceedings of IJCAI 1981*.
- Spradley, J., Mann, B., 1975. *The Cocktail Waitress: Woman’s Work in a Man’s World*. Wiley, New York.
- Stearns, P.N., Lewis, J. (Eds.), 1998. *An Emotional History of the United States*. NYU Press, New York.
- Stearns, P.N., Stearns, C.Z., 1985. Emotionology: clarifying the history of emotions and emotional standards. *The American Historical Review* 90 (4), 813–836.
- Suchman, L., 1987. *Plans and Situated Actions*. Cambridge University Press, Cambridge.
- Sundström, P., Ståhl, A., Höök, K., 2005. eMoto—affectively involving both body and mind. In: *Proceedings of CHI 2005*, pp. 2005–2008.
- Tomkins, S.S., 1995. Commentary: the ideology of research strategies. In: Demos, E.V. (Ed.), *Exploring Affect: The Selected Writings of Silvan S. Tomkins*. Cambridge University Press, Cambridge, pp. 196–203.
- Wittgenstein, L., 1953. *Philosophical Investigations*. Blackwell, Oxford.