

Emotion: Concepts and Definitions

Roddy Cowie, Naomi Sussman, and Aaron Ben-Ze'ev

Abstract This chapter deals with the task of defining and describing emotion. What do people mean when they identify emotion as a key domain for computing? How are “emotions” related to, and differentiated from, other affective phenomena? The chapter considers the definitions of emotions (and other affective states) formulated by scientists and those that are implicit in everyday language. Empirical results regarding the conception of emotion in everyday life (e.g. frequency of emotional reports in different contexts) are presented and discussed. The focus is always on the way conceptual and terminological issues impact affective computing.

1 Introduction

This chapter is about using words to describe phenomena and situations involving emotion and/or affect. The background is acute awareness, arising from a decade’s experience in technologies dealing with emotion and/or affect, that this is an area where words have a double-edged quality. They have central roles to play – in communication both between people and between people and machines and in helping researchers to order their thoughts; but they may also set traps. One of the main goals of the chapter is to help people in the area to recognise some of those traps and to deal with them.

An obvious kind of trap involves restriction. Relying on a very limited set of verbal resources can push people towards a seriously oversimplified conception of the area. A subtler kind of trap involves ambiguity. The word ‘emotion’ itself illustrates that kind of trap. Philosophers have written that emotion pervades human life (Stocker and Hegeman 1996), and in one sense of the word, it seems obviously true. That is presumably why so many people feel that it is important to take account of

R. Cowie (✉)

Department of Psychology, Queen’s University Belfast, Belfast, Northern Ireland, UK
e-mail: rodny.cowie@qub.ac.uk

emotion when they are designing artefacts, including computer systems, for humans. And yet, people also tend to agree that the word emotion ‘strictly’ refers to a very specific kind of state. States of that kind occur briefly and occasionally (perhaps once a day) in everyday life. It is a common kind of ambiguity: the word has two senses, one broad and one narrow, as do ‘cat’, ‘boat’, ‘Kleenex’, and a great many other words. Nevertheless, the consequences can be serious. Enthusiasts who argue that technology should engage with emotion are likely to be talking about something that pervades human life. Sceptics are more likely to be questioning the case for engagement with something that occurs briefly and occasionally. That kind of mutual noncomprehension leads all too easily into deadlocks that have the potential to do real damage to the field.

The need to communicate with a range of outsiders constrains the way these problems can be handled. In some areas, experts can eliminate problems rooted in semantics by agreeing among themselves to use words in specialised ways. But in this area, experts need to be able to give a fair picture of what they can do, or aspire to, to outside colleagues, funding bodies, institutions, interested firms, and others; and their systems need to be understood by, and perhaps to communicate with, ‘naïve users’. As a result, the only obvious way forward for the foreseeable future is for experts to develop a sophisticated understanding of the language in general use and the substantial issues that lie behind it. The approach in this chapter is to point people towards that kind of sophistication.

2 Plato’s Middle Ground

A key first step is to stand back from words and consider what it is that people who are interested in emotion and computing want to engage with. After that, one can ask how well the words that come to hand express the underlying interest.

Writers have often imagined a being who is as intelligent as we are, or more so, but whose mind can only process information in a strictly rational way. Probably the most famous example is Star Trek’s Mr Data. Many people who work in emotion-oriented computing seem to be motivated by a sense that a being like Mr Data would lack something, an ingredient X, that is central to being human, and that technology could and should engage more systematically with that ingredient X.

Star Trek did not invent the idea of such an ingredient. On the contrary, the Star Trek character captures the imagination because he reflects a widespread intuition. Influential versions were articulated by Augustine and, before him, by Plato.

Plato’s version (*The Republic*, Book IV) is a degree subtler than Star Trek’s. It proposes a three-part division of the mind. As in Star Trek, it was accepted that reason was something distinct. The main debate was whether the rest of mind involved one category or two. Plato argued for two. At the lowest level were pure appetites, simple and amoral, which reason either controls or is controlled by. Between appetite and reason was spirit – exemplified by anger – which is inherently attuned to social and moral issues and capable of allying with reason. On the

whole, modern research also separates off the phenomena that Plato would have called appetitive. They tend to be called ‘drives’ in modern parlance. There are exceptions, such as Rolls (1999), who takes hunger as an archetypal emotion; but that is not the norm.

Despite differences of detail, Plato and Mr Data reflect the same broad kind of intuition: that phenomena like anger are instances of something that plays a very large part in making human life what it is. ‘Plato’s middle ground’ seems an apt enough phrase to describe the domain without making specific commitments about its boundaries and contents. The point of introducing the phrase is that it seems to express what people in emotion-oriented/affective computing feel intuitively technology should engage with.

Unfortunately, translating the intuition into well-defined words and concepts is fraught with difficulty. Many words are naturally associated with the elusive category. They include emotion, feeling, expression, passion, and affect. All of them pose the same kind of problem. In the right circumstances, they can be used to designate something like Plato’s middle ground. However, each of them also has at least one other sense, corresponding to a specific part of the domain. That creates an immense potential for confusion both between parties and within a single person’s thought.

2.1 Common Terms and Their Ambiguities

When non-experts want to describe the domain as a whole, their first choice tends to be ‘emotion’ and its cognate forms (‘emotional’, ‘emotive’, etc). Negative forms in particular fit the role quite well. To say that someone is unemotional or emotionless conveys that factors which affect most people most of the time are not operating. In terms of Plato’s picture, the middle ground has shrunk to nothing, leaving the field (perhaps disturbingly) to reason and amoral appetite. Positive forms are more problematic, though.

A study for HUMAINE demonstrated the difference experimentally using video recordings (provided by a TV company) of people dealing with challenges in a novel outdoor environment. The company’s psychologist selected about 5 h of material that she regarded as representative of the types of experience found in the whole. Four raters watched the tapes and indicated moment by moment which of three categories best described their impression of the person being recorded – experiencing emotion in the full sense of the word; unemotional; or in an intermediate state involving elements of emotion, but not emotion in the full sense of the word. Ratings divided as follows. States perceived as unemotional made up 7% of the total. States perceived as emotional in the strong sense made up 14%. The remaining 79% was perceived as intermediate, with elements of emotionality, but not emotion in the strong sense.

Studies like this can only give ball park estimates, but those are all that matters here. It seems clear that one sense of the word emotion refers to something that

makes up a small part of human life – perhaps about a sixth even in challenging circumstances. But the word has another sense, which refers to something much less sharply defined, and much commoner. That sense comes to the fore when the negative form, ‘unemotional’, is used. There is an intriguing variant of the same point in Augustine (1984): he claimed that complete absence of emotion (*apatheia*) did not belong in this life (City of God, XIV9).

There is not the same direct evidence for other words, but related patterns seem to hold. When the heroine in a romantic novel sobs that the hero is utterly devoid of feeling, the void that she has in mind probably corresponds quite well to Plato’s middle ground. However, ‘feeling’ in its more precise sense conveys a domain that is different from Plato’s in important ways. First, it refers strictly to phenomena that are subjective and part of consciousness. Second, it very definitely includes phenomena that are much more basic than the ones Plato had in mind, as in ‘I have no feeling in this leg’.

‘Expression’ is included here mainly because respected figures argue that what is currently described as research on emotion should talk about expression instead (Campbell, 2003). However, ‘expression’ in the precise sense misses the mark in much the same way as ‘feeling’, but in the opposite direction. First, it refers strictly to phenomena that are objective and observable. What lies beneath (such as feelings) is strictly no part of the domain – even if it is being expressed, but especially if it is not. Second, in the strict sense, even some objective signs are excluded – a racing heart, rash behaviour at the steering wheel, selective attention, and so on.

The term ‘passion’ is included for a similar reason. At one stage, it was the word that philosophers most often used to describe phenomena that we would call emotional. In the present era, though, it tends to have a narrower sense, referring to states where feeling overwhelms reason. Some of Plato’s middle ground is like that, but not very much.

‘Affect’ is a word that deserves special attention, because it is much used in the area, and it has a very curious semantic profile. It is rarely used in everyday discourse. Insofar as it has a generally accepted meaning, it signifies something akin to emotion, but broader in some sense. Experts have taken it up and given it a great variety of more precise senses, often grounded in a theory which implies that emotional and emotion-related phenomena divide naturally in particular ways.

The term was taken up early in the history of psychology. William James noted that German writers used it to refer to “a general seizure of excitement (. . .) which is what I have all along meant by an emotion” (1920, p. 358). Freud gave it a specific sense in the context of his psychoanalytic theory, and that sense gained some currency (Rapaport, 1953). Medicine also adopted the term, but with quite inconsistent usages. For instance, Dark defines it explicitly as an inward state: ‘The feeling-tone accompaniment of an idea or mental representation. It is the most direct psychic derivative of instinct and the psychic representative of the various bodily changes by means of which instincts manifest themselves’. In contrast, Abess defines it as ‘observable behavior that represents the expression of a subjectively experienced feeling state (emotion)’.

In the mid-twentieth century, psychologists like Hilgard (1980) and Sivan Tomkins took up a philosophical tradition of using ‘affect’ as a name for a division based on Augustine’s. Tomkins described affect as the person’s ‘heart, his feelings, his affects’ (1964, p vii). His usage was very broad indeed. He explicitly included overt signs as part of affect and used the term to cover not only the standard emotions but also arousal, hunger, pain, commitment, and various other states that would not generally be called emotional, some of which Plato would clearly have excluded from his middle category.

Another set of usages has gained currency more recently. Panksepp (2003) describes affect as ‘the feelings associated with emotional processes’, and Russell (2003) describes it as similar to ‘what is commonly called a feeling’. Both tie affect specifically to feeling. However, they also make it clear that their sense of the word has another level. They use ‘affect’ to describe states with a dual character, which involves both experience and physiology. For instance, Russell (op. cit.) describes core affect as ‘a neurophysiological state that is consciously accessible as a simple, nonreflective feeling’. When the activity of the neural systems involved is reflected in consciousness, we experience it as feelings; but the systems can be active without being reflected as feelings.

The appeal to physiology has hidden subtleties. There is a recurrent suggestion that a person’s affective state could be established by monitoring the relevant neural events, using current or near-future technology. In contrast, nobody expects to be able to decipher a person’s beliefs by similar means in the foreseeable future – presumably because correspondences between beliefs and neural events are felt to be much more intricate. A much grosser correspondence seems to be assumed in the case of affect as understood by Panksepp, Russell, and others. That may hold for some aspects of anger, for instance, but correlates of the fact that anger is directed towards a particular person or thing, and is felt to be morally right, seem likely to be as subtle as the correlates of belief.

Research that describes itself as ‘affective computing’ seems on the whole to lean towards a use of the term that is broadly similar to Panksepp’s and Russell’s. It is likely (though not certain) to be particularly concerned with states that might be identified by monitoring some relevant neural events. If so, it represents a particular way of approaching Plato’s middle ground, guided by a particular scientific model of what gives that ground its character. That is why this chapter does not use the term ‘affective computing’ to describe the whole enterprise of trying to engage computing with Plato’s middle ground. It would be like using the term ‘Catholic’ to describe all mainstream Christian denominations – defensible in principle, but likely to be confusing in practice.

The particular ambiguities that have been discussed involve principles that have quite wide-ranging effects. There is a well-known phrase *pars pro toto*, meaning ‘the part stands for the whole’. It captures a feature of discourse in the area, which is that at least in casual conversation or writing, terms like ‘feeling’ or ‘expression’, or simply examples like anger, can stand well enough for the whole of Plato’s middle ground. In formal contexts, though, an opposite principle seems to apply: *pars invadet totum*, the part usurps the whole. An investigator may begin using a term

like ‘expressiveness’ intending to refer to Plato’s middle ground or at least most of it – *pars pro toto*. But as a research effort develops, the narrow sense tends to shoulder others aside, so that it becomes difficult to justify considering anything that is not expression in the narrow sense and mandatory to consider aspects of expression that are not particularly relevant to the original conception – *pars invadet totum*.

The way words move between broad and narrow senses causes a multitude of difficulties; and yet it is not something that can be eradicated from language, not least because it is useful to non-experts. However, there are strategies that make it easier to avoid the most negative consequences.

2.2 Systematising Vocabulary/Emotion Terms

A time-honoured strategy for dealing with ambiguity is to set key terms in phrases that direct people reasonably reliably to one sense rather than the other. A key attraction of the strategy is that it gives people access to multiple ways of bounding and dividing the domains associated with Plato’s middle ground – divisions which everyday language implies, but does not make it easy to separate cleanly. Various options of that kind have been explored within HUMAINE.

‘Pervasive emotion’ emerges as a reasonably satisfactory way to refer to whatever is present in most of life, but absent when people are emotionless (which the data given earlier suggest happens rather rarely). The term is adapted from Stocker and Hegeman (1996). It is the single most convenient description of a domain roughly coextensive with Plato’s middle ground. It is surprisingly difficult to find a term that expresses the narrower sense satisfactorily. ‘Emergent emotion’ poses fewer problems than other options that have been considered. It reflects a widely held interpretation of this kind of state, which is that it involves multiple elements coming together to form a distinctive *Gestalt* (Scherer and others describe the effect as synchronisation) which either dominates the way a person acts and thinks or needs to be held in check by a deliberate effort. To complete the set, ‘emotional life’ has been used to refer to the sum total of the states, processes, experiences, and actions that are substantially influenced by pervasive emotion and therefore distinguish human life as it normally is from the life of a being who is always and completely emotionless.

Both feelings (which are internal) and expression (which is public) are key elements of emotional life. There are feelings that we do not consider emotional (such as pain). It is not self-evident why people distinguish them from emotional feelings. Plato suggests a distinction based on different relationships to moral and intellectual systems; Ortony and his collaborators propose that the hallmark of emotional feelings is an element of positive or negative evaluation (Ortony and Turner 1990). ‘Passion’ is used technically by philosophers to distinguish emotional phenomena that carry a distinctive kind of compulsion. In the context of computing, the natural default is to assume that ‘affect’ refers to phenomena, including feelings, which

(*ex hypothesi*) arise rather directly from the activity of particular neural systems. More specific senses can be identified by a standard phrase (e.g. ‘core affect’) or the name of someone who uses the sense (affect in Panksepp’s sense or Freud’s sense, etc.).

The point of systematising vocabulary like this is to help people to avoid traps. There are various kinds of traps that may be easier to identify and avoid once the distinctions are registered. An obvious example is setting out to take account of what intuitively seem to be pervasive human characteristics; using the term emotion to describe them (as is natural) and being drawn (by the principle of *pars invadet totum*) into research on the rare phenomenon of emergent emotion (for instance, collecting databases full of examples of emergent emotion or building agents that simulate it). Similarly, if the term used to describe the pervasive characteristics is ‘affect’, it is natural to be drawn towards a search for correlates of hypothetical biological processes and to lose sight of the cognitive and moral aspects that mark off the pervasive phenomenon from phenomena like thirst. Choosing to study feeling (which is private in the narrow sense) makes it easy to lose sight of the interpersonal aspects of the domain; choosing to study expression carries the temptation to gloss over hard questions about what lies behind the expression and what is read into it; and so on. Many of these carry with them the trap of conveying to the general public that one is going to do much more than one actually is – a trap that can be deadly in the long run for the reputation of disciplines that fall into it.

Traps like these are nobody’s fault. They exist because emotional life is a huge, complicated domain, and people are continually looking for ways to make their dealings with it manageable. Proposed solutions that ignore those realities are unlikely to work.

2.3 Coda

An obvious question should be faced before the end of this section. Plato has provided the framework for the discussion: but why should anyone take his framework seriously?

Tripartite divisions have been surprisingly widely accepted. A division derived from Plato via Augustine, using the terms cognition, conation, and affect, remains widely used (Hilgard 1980). Ortony and his group (Norman et al. 2003) propose another related division (using the terms reflective, affective, and reflexive). However, the root answer is that what matters here is not whether Plato’s analysis is correct, but whether it reflects ideas that people bring to the field, and thereby helps them to clarify their thinking. A decade in the field suggests that people who enter it often do have something like Plato’s middle ground in mind. Other schemes involving slightly more or slightly different subdivisions may capture that intuitive sense marginally better, but the differences are not particularly important in this context.

There may be people whose motivation is different – whose aim is simply to study a domain that is fully and accurately captured by the strict sense of one of the other terms – emergent emotion, or affect in the sense of Panksepp or Russell, or

passion. They may be quite clear in themselves about what they want to study. But if they want to convey it to others, even they may benefit from having some sense of what another person might mistakenly assume – *pars pro toto* – that they were studying.

On the other hand, people who are drawn to something like Plato's middle ground intuit that what they are trying to engage with makes up such a large part of human life that its importance for technology hardly needs arguing. Evidence like the TV study cited above indicates that their intuition is right, however tricky it may be to put it into words.

3 Describing Fragments of Emotional Life

This section shifts focus from the macro-task of naming a large domain to the micro-task of describing individual parts. It tries to draw together the main descriptive resources – both terms and concepts – that are needed to convey what is happening in a particular situation where emotion is a key factor. Throughout the section, 'emotion' is used in the sense of 'pervasive emotion' (*pars pro toto*) unless the context indicates otherwise.

Between them, philosophy and psychology provide a very rich set of descriptive resources. It is easy to underestimate the resource, because key ideas are often associated with (apparently) conflicting theories, as if they were alternatives. Broadly speaking, conflicts tend to involve claims that a particular set of concepts captures the central essence of emotion. It is not clear how much technology needs to be concerned with claims about essence. If those debates are set aside, contributions by a range of theorists can be seen as acute descriptions of factors that may or may not define the essence of emotion, but that are certainly relevant to describing what one can expect to find happening when emotion is present.

The sheer number of factors that has been identified is a key point in itself. It is a sharp reminder that like it or not, the domain of emotional life is massively complex; and that there are good reasons to be wary of any model that appears to reduce it to a few simple concepts. Issues are grouped under a few headings here for convenience. Different headings could certainly be chosen, but these provide a structure and progression that seem useful.

3.1 Units

Emotional life is typically divided into two major types of part, illustrated by William James's urgent reaction to a bear in the woods, on the one hand, and by Lord Jim's lifelong shame at jumping ship as a young man, on the other. The term 'emotional episode' will be used for a case like James's, where the person's mental state changes briefly but deeply, and 'established emotion' to describe units like Lord Jim's shame, which are quite likely to last for a lifetime. An established emotion is likely to underlie many emotional episodes, but it tends to be dispositional most of the time (Goldie, 2000).

These large-scale units in turn are individuated by features of various different types. Some are components (that is to say, processes or structures in their own right); others are attributes (that is to say, properties of processes or structures). The rest of this section is concerned with setting out the key types of component and property that give an individual emotional episode or established emotion its character.

3.2 *Dimensions*

It is often natural and useful to describe emotion in terms of a few dimensions (it may or may not be theoretically profound). Three dimensions are very standard. *Valence* describes the value (positive or negative) of the feelings involved. *Activation* describes the strength of the individual's disposition to act. *Potency* describes the individual's sense that he/she has the power to deal with relevant events.

Many other dimensions have been proposed, of which two will be mentioned here. Unpredictability is included because it emerges as a key factor in a particularly well-constructed study (Roesch et al., 2006). *Engagement* (as opposed to detachment) is rarely mentioned in the psychological literature, though Ortony (2002) used the term 'caring' to express what seems to be a related concept. However, it does concern technologists working on 'presence' in virtual reality (i.e. the sense of being materially engaged with the virtual surroundings, rather than essentially distanced from them). In that context, it is widely assumed that full emotional responsiveness to virtual surroundings implies, and depends on, engagement with them (Huang and Alessi, 1999).

3.3 *Feeling*

Distinctive kinds of feeling are among the obvious hallmarks of an emotional episode. They are notoriously difficult to describe. Two main ideas about description are standard. One involves dimensions. Russell and his collaborators refer to the characteristic feeling element of emotion as 'core affect', implying that it bears a special relationship to neurophysiology. They suggest that it is characterised primarily by valence, secondarily by activation (Russell and Feldman Barrett, 1999). On the other hand, William James (1884) proposed that the feeling element of emotion consisted of awareness of somatic changes (in heart rate, breathing, etc.) associated with the emotion. His idea seems to be partly true: injuries that prevent detection of somatic changes do alter the quality of emotional feelings, but they do not eliminate them.

Most investigators accept that emotion can exist without feeling, most obviously because the other hallmarks of an emotional episode can exist without conscious emotional feelings. (Note that in James's example of the bear, fear reactions precede fear feeling.) That is one of the arguments for defining emotion in terms of the activity of systems with a particular link to feeling rather than feelings per se. According to authors like Panksepp (2003), the work of these systems need not

impinge on awareness; but when it does, it has a characteristic quality, which we convey by saying that it is felt rather than analytic. *Ex hypothesi*, these systems generate evaluations that are felt rather than calculated, inclinations to act that are felt rather than deliberately decided, and so on. It would help to explain the lasting appeal of James' proposal if they had close biological links to the systems that generate visceral feelings. Its strength would then come from an intuitive sense that the same kinds of system were in play in emotional and visceral feelings – that the two were cut from the same cloth, so to speak.

3.4 Appraisal

One of the hallmarks of an emotional episode is nicely captured by Ben-Ze'ev's phrase 'partial perception' (2000): it involves a selective grasp of a situation, which highlights what is relevant to the 'weal or woe' (Arnold 1960) of key players. The best known developments of that idea propose relationships between emotion categories and value-oriented 'appraisals' of the situation.

There are many specific descriptions of appraisal. A well-developed example is due to Scherer's group (Sander et al., 2005). It describes a sequence of 'stimulus evaluation checks' which makes sense logically and fits data collected by the group. It proposes that the onset of emergent emotion involves a series of checks, in the following sequence:

- Relevance (including sub-checks for novelty, intrinsic pleasantness, and relevance to the subject's goals and needs),
- Implications (including sub-checks for causal attribution, outcome probability, discrepancy from expectations, goal conduciveness, and urgency),
- Coping potential (including sub-checks for the controllability of the event and the subject's power to affect its course and/or to adjust to its consequences), and
- Normative significance (including sub-checks concerned with the way outcomes relate to one's own values and to society's).

One of the analyses that has had most impact in technology, the framework proposed by Ortony et al. (1988), is also rooted in appraisal theory. Their analysis in turn exists in both an extended and a reduced version. It is beyond the scope of this chapter to evaluate different forms of appraisal theory.

3.5 Emotional Colouring

It is natural to picture the result of appraisal checks as a representation in which descriptors specifying emotion-related qualities are attached to significant things and relationships in the relevant situation. The idea can be expressed by saying that the representation is emotionally coloured – with the colours indicating whether key features of the surroundings are pleasant, conducive to the person's goals, within the person's power to control, morally acceptable, and so on.

The metaphor extends to the general meaning of concepts as well as impressions of individual situations. A large body of work by Osgood and his collaborators (Osgood, 1957) showed that the emotional colouring of concepts (their term was ‘feeling tone’) could be summarised reasonably well in terms of three dimensions (evaluation, potency, and activity). It would be interesting to revisit the area in the light of the richer descriptive systems that appraisal theorists have developed since.

3.6 Action Tendency

There is a long-standing recognition that emotion tends to close the gap between having an impression of the situation and acting on the impression. For instance, Aristotle cited the impulses to/for revenge as a defining feature of anger. Frijda (1987) reintroduced a related concept in the modern era and his term ‘action tendencies’ is widely used. He argued that tendencies to act in particular (biologically significant) ways were integral to emotion and were central to distinguishing among emotions with a direct biological significance – tendency to approach is the kernel of desire, tendency to avoid is the kernel of fear, tendency to reject is the kernel of disgust, and so on.

There is another sense in which emotion is bound up with instigating action, which may or may not be fundamentally separate. It involves motivation. It is not clear how tight the connection is. Motivation is often linked more to Plato’s lowest category of appetite (hunger, thirst, pain). However, emotional colouring is obviously a factor in motivation (you will do more for someone you like or fear); and emotional episodes certainly motivate or demotivate. Sylvan Tomkins (1991) proposed that emotions act as amplifiers that modulate basic drives. It is not unlike a metaphor that Plato used: he imagined emotion as the good horse that responded directly to the driver (rationality), drawing its less co-operative companion (appetite) along with it.

3.7 Expression

Actions with a communicative element are among the most characteristic components of emotional episodes – smiling, weeping, screaming, and so on. Theorists from different traditions have understood these in substantially different ways.

Simplifying grossly, accounts that appeal to evolution have tended to assume that expressions of emotion are produced by innate mechanisms which automatically generate external signs of significant internal states, with socially defined operations (display rules) capable of concealing or mimicking the innate patterns (though not usually perfectly). In contrast, social psychologists argue that the patterns are fundamentally communicative: smiles are directed to people, not automatic externalisations of an inner state. The two lead to different research strategies. For instance, evolutionists assume that expression which is uncontaminated by display rules has a privileged status and should be sought out. For social psychologists, it is

a fiction that draws people into studying situations that are ecologically unrepresentative (misguided in rather the same way as trying to study sitting behaviour without the complicating factor of a chair).

3.8 Emotional Modes of Action and Cognition

Emotion affects not only what people do, but also the way they do it (of course, the line is often blurred). Some of the effects flow from underlying shifts in the way people perceive and think under the influence of emotion. There are well-documented examples at many levels of cognition.

A practically important example of effects on attention is called ‘weapon focus’ – exclusive concentration on a single, focal detail of a scene (the gun) to the exclusion of other features which are actually important (the gunman’s face). The effect seems to be due to the perceptual processes that evoke emotional responses (Laney et al., 2004). There is a substantial literature on the way anxiety affects perceptual and related process – attentional control, depth of processing, and speed of processing. Eysenck’s work on anxiety illustrates a well-developed analysis of the issues surrounding these effects (1997, 2007). Positive conditions tend to generate extensive and well-organised memories, and positive affect promotes their recall later. It also fosters flexible and creative thought, can speed decision making, and affects risk taking, not necessarily in obvious ways (Isen, 1998). Negative moods tend to increase people’s impression of the effort that a task requires (Gendolla and Krüsken, 2002). Marked emotion tends to reduce coherent verbal communication (Cowie and Cornelius, 2003).

These effects are practically important for emotion-oriented computing. Consider, for example, the implications of ability to recognise when emotion is impairing a driver’s perception of risk, or a pupil’s ability to learn, or a manager’s ability to communicate clearly, or a worker’s readiness to sustain effort, and so on.

3.9 Connectedness

Usually (perhaps always) describing an emotional episode depends on referring beyond the person who has an emotion experience to various significant objects and significant others. That is already implicit in several of the points above, but deserves to be drawn out. An appraisal is an appraisal *of* people, events, or things; and expressions of emotion tend to be directed *to* particular people in the context of an audience (couples know how dramatically the sudden appearance of a child or an in-law can affect the expression of various emotions).

Philosophers use the term intentionality to express the fact that emotions are *about* something, which is called the object of the emotion (as against the subject who experiences it). ‘Connectedness’ is broader and aims to cover both that kind of

linkage and the linkages involving others who are involved in an emotional episode, but not the object of the emotion.

The list of connections above is far from exhaustive. An emotional episode may be about one thing (mortality, for instance); prompted by another (a poetry reading, perhaps); with causal roots in events long past (such as a bereavement). Shame before an audience which approves of an action may arise because another audience (present in the mind rather than in reality) would not approve. There is no obvious end to the permutations. In addition, mixed emotions often involve connections with multiple different events, bearing different emotional colourings – gladness that a gap in life is to be filled, sadness about the loss that created the gap, concern that it might still not work out, and so forth.

These issues have points in common with claims about ‘groundedness’. Advocates of ‘groundedness’ argue that certain kinds of representation can only emerge by a causal process from the reality that the agent is part of. However, it seems fair to say that their concern is with a different level of connection, defined by the fact that a history of causal interactions has moulded the symbolic medium in which connections, with current or past events, are expressed. It is not clear what hinges on the existence of such a level, and the two concepts are probably best kept apart.

3.10 Impressions of Emotion

Emotional episodes typically involve more than one person; and when they do, understanding how signs of emotion are registered is as essential as understanding how they are emitted. Scherer (2003) recently reasserted the point in a neat form using a lens model adapted from Brunswik. There are several ways of conceptualising registration.

Detection paradigms consider whether an objectively verifiable state is identified correctly. That approach is clearly appropriate in certain application areas – emotional intelligence tests, for instance, and lie detection (where emotion-related signs are assumed to be pivotal).

Experiential paradigms are broadly comparable to certain areas of psychophysics, where it is accepted that subjective experience may have its own dimensions. In a parallel way, it makes sense to consider whether (for instance) dimensional descriptions capture the way we perceive other people’s emotions under certain circumstances, whether or not they capture the intrinsic nature of our own and so on.

Control paradigms consider how variables affect behaviour rather than experience. The two can be very different. For instance, it is well known in perception that the behaviour of visually guided grasping is not affected by variables that distort conscious reports of size and distance. There is evidence that in a parallel way, responses to others’ emotion may reflect variables that are not reflected in conscious impressions of the other.

Timing needs to be considered alongside these distinctions. Taking time to identify a single static state is not the same as registering in real time how a person's emotional balance and focus is shifting, which is what people have to do when they participate in an emotionally coloured interaction. Perceived flow of emotion seems an apt term for the kind of impression that underpins real-time interaction.

Subjects perceive (or fail to perceive) their own emotions as well as other people's. That would seem to involve forming explicit representations of changing flows and pressures that are at work in their own heads (so to speak). Helping people to perceive their own emotions is one of the application areas that is regularly considered for emotion-oriented computing.

3.11 Category Labels

This section has listed various kinds of resources that are relevant to describing a particular fragment of emotional life. It has deliberately left until last the resource that people typically consider first, that is, words like 'mood', 'anger'. The reason is that people find it very easy to think of emotional life as a collection of events that correspond reasonably closely to salient category labels. It is a prime case of *pars invadet toto* – attention is pulled onto the special cases which are close to category archetypes, leaving the mass of everyday phenomena that are far from the archetypes sidelined.

Databases that use naturalistic material highlight the issue. Research teams repeatedly observe that what they find often is not particularly well described by any standard category label and often seems to involve multiple categories (Douglas-Cowie et al., 2003, 2005). Even when category labels do fit, they do not in themselves provide information that is crucial to understanding the events (for example, triumph at a football match is likely to differ in a great many ways, visible and invisible, from triumph in a court room). The point is not that category labels have no part to play. On the contrary, they are considered at length in the next section. The point is that they are one kind of resource among many.

A significant divergence from that view should be acknowledged. Many theorists argue that a few qualitatively distinct neural systems give rise to the whole of emotional life, and that the most important category terms are linked very directly to those systems. The term 'basic emotion' is used in various ways, but one of them is to describe states that are hypothesised to relate simply and directly to one of these systems.

It is not obvious how important that idea is in practice for emotion-oriented computing. It may be that a few systems underlie the complex emotions involved in attending the funeral of a former colleague who had been able but quarrelsome; but it does not necessarily do much more to illuminate the emotions than knowing about retinal receptor types does to illuminate the complex visual experience of watching faces at the funeral.

In that situation, technologists may make different choices. Some will judge that it is worth pursuing the idea that there are a few basic systems underlying emotional life. Others will choose to treat category labels as simply one resource among many

to be used in engaging with an inherently complex set of phenomena. Discovering the relative success of the two approaches in computing may contribute a good deal to resolving the theoretical disputes within psychology: it is hard to see how anything but a computational approach can establish how much of the complexity of real emotional life each can accommodate.

The kind of framework that has been set out in this section is intimately related to the development of databases. Many of the ideas were prompted by the surprising difficulty of describing the phenomena that were observed in naturalistic recordings. In turn, the descriptive system in the HUMAINE database is a practical simplification of the ideas presented in this section.

4 Classifying Emotion-Related States

One of the most natural ways of thinking about emotional life is in terms of something like a taxonomy of states. If a standard emotion word such as ‘fear’ refers to a species of individual states, then ‘emergent emotion’ refers to a genus, and ‘emotional life’ refers to a family of states. The idea is not as straightforward as it might look, but it does provide a useful framework in which to set out the various kinds of description that deal with states related in some way to pervasive emotion.

4.1 *Mindsets and Personal States*

People are sceptical, and reasonably so, when emotion and related terms are used as a catchall for everything that is conveyed by non-verbal communication. Clearly states which can reasonably be thought of as related to emotion are only part of larger domains. A natural name for a domain one step larger is ‘mindsets’. Included at that level might be social states (dominance, deference, and so on) and cognitive states (confusion, interest, etc). ‘Personal states’ is a natural term for a broader domain still, including, for instance, states of health and well-being (ill, vigorous, and so on) as well as mindsets.

It is a useful approximation to say that emotion-related states are a particular kind of mindset. That allows (for instance) people who are working on classification of states observed in meetings to say that they are interested in states which are mindsets, but not (in the main) emotional. It is approximate rather than exact because an actual person (unlike Mr Data) is unlikely to be in a state that has no emotional elements. The division is between descriptors which do and do not refer to emotion-related factors in the person’s state, not between states that do and do not have emotional elements.

To complicate matters still further, a large proportion of terms that are not explicitly emotional nevertheless include emotion as a likely factor. For example, courage seems to be essentially a behaviourally defined state; and yet its meaning has intimate links to emotional factors (such as controlling fear). Is courage, then, to be considered an emotion-related state? Lazarus (1999) illustrates a slightly different

kind of connection when he observes that it is unrealistic to discuss stress without reference to emotion, even though they can be separated conceptually.

Questions like these probably have to be handled pragmatically and in context. So if one wants to model a brave man's behaviour, it may be enough to consider courage as a matter of risk assessment; but instilling courage probably depends on engaging with its emotional aspects.

4.2 *Generic Emotion-Related States*

On the taxonomic metaphor, emotion-related states make up a family and can be divided into genera – obvious examples being moods and emergent emotion. One might assume that there would be well-accepted ways of subdividing emotion-related states into genera and that they would be backed up by data on prevalence. It is clearly of interest to emotion-oriented computing to know which kinds of states are actually common, since (other things being equal) it makes sense to orient systems to common phenomena rather than rare ones.

In fact, there have been few systematic attempts to provide a set of categories that between them cover the whole domain of emotion-related states. A table due to Scherer, given below, provided a starting point for research on the topic in HUMAINE. The research extended the table in two stages: the first a priori and the second empirical.

The a priori stage started by considering features that distinguish emotion-related states from other mindsets. The most obvious of those are dimensions of emotion that were listed in the last section: valence, activation, potency, and engagement. It makes sense that the genera might include states that are distinguished by an unusual level of one factor, but not the others. Mood falls neatly into that framework, since it is often described as a state distinguished mainly by valence. Intuitively, there do seem to be states that correspond to the other dimensions in similar ways, involving heightened or lowered sense of potency or control, activity, and engagement or seriousness about events at the focus of attention.

Similarly, the states listed in Scherer's table are either stable or follow a relatively set trajectory. One type of state that changes more freely with time has already been

Table 1 Genera of emotion-related states (after Scherer et al., 2004, p 11)

Design features	Emergent emotions	Interpersonal stances	Moods	Attitudes	Affective dispositions
Impact on behaviour	++	+		+	+
Intensity	++	+	+	+	
Rapidity of change	++	++	+		
Brevity	++	+	+		
Event focus	++	+			
Appraisal elicitation	++				
Synchronisation	++				

mentioned, that is, established emotion, which is usually latent, but occasionally translates into an emotional episode. On a shorter time scale, there can be oscillation between a sustained mood-like state and outbursts of emergent emotion. Recognising transition as an issue suggests that transitional states are sometimes sustained: people can simmer on the edge of anger, but not quite succumb. The category ‘emotionless’ was also added, for obvious reasons.

Empirical studies followed up those ideas. Naïve participants were given lists of genera based on the reasoning outlined above. For each category, they were asked to assess whether they had experienced episodes that fitted the description; and if so, to give brief accounts of them. Most accounts fitted the categories and confirmed that they were recognisable, but some described experiences that were not well captured by the a priori framework. That led to additions in two areas. Stances towards things and situations were added to stances towards people; and a category involving more enduring orientations to people (described as bonds) was also added. At the same time, the term ‘attitude’ was abandoned, because participants clearly used it to mean something quite different from what was intended. It is noticeable that the term has also acquired quite different meanings in different academic literatures, taking on one technical sense in a literature derived from Ajzen and his colleagues (1988) and another in linguistics (O’Connor, 1973), which has sometimes been equated with Scherer’s term ‘interpersonal stance’ (Wichmann, 2002). Like ‘affect’, ‘attitude’ seems to have a facility for picking up multiple meanings; and it needs to be used cautiously for that reason.

The revised list was used in an ‘ambulatory study’ (Wilhelm et al., 2004). Ten participants were given a protocol in which each generic category was named, described, and illustrated with an example given by a participant in the previous study. Each of them was then contacted by phone 50 times at random times over a period of weeks. They responded by identifying the generic descriptor that best reflected their state at the time. The main results are summarised below:

Established emotion	0.9%
Emergent emotion (suppressed)	1.7%
Emergent emotion (full-blown)	1.5%
Mood/emergent emotion oscillation	1.5%
Mood	36.1%
Stance towards object/situation	25.6%
Interpersonal stances	2.4%
Interpersonal bonds	4.1%
Altered state of arousal	21.9%
Altered state of control	3.9%
Altered state of seriousness	0.4%
Emotionless	0.0%
None of the above	0.0%

These data reinforce points made earlier through the TV study: people are rarely emotionless and not often in a state of full-blown emotion. They also show that between the two is a variety of states which are very common – some of which have established names, but by no means all. Hence if emotion-oriented computing wants to address the emotion-related states which occur commonly, it cannot let itself be guided exclusively by the labels that are available in everyday language.

The work described here clearly needs to be developed. But it provides at least a preliminary overview of the states that make up Plato's middle ground and some protection against being drawn (*pars invadet toto*) into studying emotion in a sense that accounts for a very small proportion of emotional life.

4.3 Specific Emotion Terms

Words like anger, joy – what is being called the species level here – clearly have a kind of priority in the domain of emotion. In Rosch's terms, they appear to be the basic level in this domain – as 'dog' or 'cat' do in the domain of animals. It is natural to assume that they are essentially the names of states and more specifically of affective states. A good deal of research has gone into showing why that kind of assumption is at best a first approximation and can be quite seriously inappropriate for emotion-oriented computing in particular.

There are sources which offer to analyse words at this level as co-ordinates in a space with a low number of dimensions. Whissell's dictionary of affect is a particularly thorough example of that approach, and it is reasonable to interpret the co-ordinates as a description of affect. That kind of analysis is often useful, but it is very far from sufficient.

A first kind of problem concerns variation in the scope of words. Words like anger and love can famously refer to many different states – hot and cold anger, sexual and nurturant love, and so on (Russell and Fehr, 1994; Sternberg, 1988). They can also refer to one-off emergent episodes (anger at a rude shop assistant) or established emotion (anger at UK policy in Iraq). Linked to that, there is evidence that there are material differences between the anger evoked by things that are physically present and the anger evoked by remembering past events (Stemmler et al., 2001). Observing naturalistic data also underlines the frequency of states which do not exactly fit any category, either because they are intermediate or because they are blends (using that term to mean that two or more emotions seem to be coexisting – happiness at one aspect of the current situation, sadness at another). If the word 'anger' describes a species, it is a species much more variable than (for instance) dogs.

A deeper problem is that specific emotion words do not simply refer to inner affective states (however varied). Their meaning is bound up with a variety of complex judgments, many related to the fact that words in everyday language must allow people other than the subject to apply them (otherwise they could not form part of a common language). Because of that, words that involve inner states must be rather complicated instruments whose rules of use are logically bound to both intra- and interpersonal elements.

The point is taken up in an article by Sabini and Silver (2005) and a reply by Cowie (2005). Sabini and Silver argue that terms like jealousy and anger, shame and embarrassment may refer to the same affective state. The difference lies in factors surrounding that affective kernel, some internal to the subject experiencing the emotion, some external. The reply draws out the implications for the task of assigning everyday emotion-related words as humans do and identifies eight types of consideration that appear relevant to the task. Clearly one type of consideration involves the internal feeling state. But assigning emotion terms also depends on the objective events which prompt an emotion (if a successful person has genuinely insulted an unsuccessful one, the latter's emotional response would be called anger; if the latter has perceived an insult in quite innocent behaviour, it would be called envy). It depends too on evaluation of the person's character (the word 'shame' is applied when a person accepts that negatively evaluated actions reflect a genuine deficiency in themselves, 'embarrassment' when he or she does not). Similar points can be made about the other types of consideration: these involve the person's appraisal of circumstances; the involuntary signs that he or she gives; his or her choice of action; the manner in which the action is undertaken; and the observer's evaluation of action.

Obviously these ideas are closely related to what has already been said about describing fragments of emotional life. What they add is that assigning standard emotion words is not a simpler task that can be dissociated from the complexities described there. On the contrary, the conditions for using specific emotion terms are bound up with the overall complexity of emotional life.

Related but distinctive problems arise over the right to use particular words. A parent may describe a child as sulking. A machine which attached the same label to the child would be presuming a right to make moral judgments which the recipient might dispute and might expect to be smashed. In general, it is not at all obvious what rights people might attribute to computers; and that means there are open questions about the emotion words they could or should use.

These considerations are not abstruse entertainment for the philosophically minded. If emotion-oriented systems are to use emotion terms appropriately, they need to use them in accordance with the complex criteria that actually govern their use in natural languages. That means recognising that although it is a convenient approximation to think of specific emotion words as labels for species of states, it is nevertheless an approximation.

5 Where Have All the Theories Gone?

This review has said very little about theorists and their positions per se, though it has drawn ideas from them liberally. This section touches briefly on major theoretical positions for the sake of reference.

Relationships between emotion and the body have been a recurring issue. Descartes depicted passion as the mind being disrupted by turbulence in the body. William James identified emotional feelings with awareness of bodily changes. His

son-in-law and antagonist, Walter Cannon, identified them with changes in the mid-brain. Late twentieth century research has used brain imaging to identify brain centres associated with emotion, led by figures like Damasio, LeDoux, and Davidson.

Darwin can be seen as a branch from that line. His view of emotion was biological, but emphasised the evolutionary constraints that made particular behavioural patterns adaptive, in humans and animals. Frijda (1987) developed a distinctive evolutionary approach which emphasised the action tendencies associated with emotions.

An influential synthesis of these approaches, most strongly associated with Ekman, uses the concept of basic emotions. By that he means that emotions are of several discrete types, each with a cluster of characteristics, laid down by evolution and rooted in discrete physiological systems. A great variety of approaches are broadly similar, for instance, those of Plutchik (1980), Cosmides and Tooby (2000), and many others. Some of these emphasise the discreteness of the hypothetical affective systems, others their ability to interact with cognition (e.g. Panksepp, 2003).

Cognitive approaches emphasise the integral part that cognitive processes, and particularly appraisal, play in emotion. They are usually traced to Arnold (1960), with influential formulations due to Lazarus, Le Doux, Ortony and his colleagues, and Scherer. Strong versions of cognitivism regard emotions as informationally encapsulated brain processes (LeDoux, 2000), whose feeling component is relatively unimportant. Weaker versions see emotion as essentially an amalgam of cognition and motivation (Lazarus, 1999).

In contrast to evolutionists, social constructivists emphasise the role of culture in giving emotions their meaning and coherence (e.g. Averill, 1980; Harre, 1986). The position held by Russell (2003) can be regarded as a distinctive variant, which considers affect an underlying biological substrate that enters into a variety of processes shaped by social and other factors. Other social theorists have argued that emotion is fundamentally an attribute of interactions between people rather than of individuals (Parkinson, 1995).

Even this short summary gives some sense of the variety of theories that carry weight in the field. That reflects the fact that nobody has yet identified a single, unifying kernel round which all that is known about emotion can be organised in a completely coherent, satisfying way. However, it does not reflect a field in utter turmoil. A large body of knowledge exists, and the bulk of the chapter has tried to reflect it. It remains a challenge to find a truly satisfying way of organising it. The scale of the challenge should not be underestimated.

6 Conclusion

From top to bottom, emotion language is more complex than it looks. That gives rise to traps when people forge ahead relying on a model which is or seems appealingly simple, but which in fact conceals both the complexity of language and the complexity of the real phenomena involved.

The strategy of this chapter has been to alert people to the complexities at both levels. No doubt many readers will find that thoroughly unsatisfying and look for articles that offer more concrete or elegant prescriptions. It may be that when they have worked their way through a sufficient number of traps, they will come back.

References

- Ajzen I (1988) Attitudes, personality, and behavior. Open University Press, Milton Keynes
- Arnold MB (1960) Emotion and personality. Columbia University Press, New York, NY
- Augustine (1984) The city of God (trans: Bettenson H.). Penguin Classics, London
- Averill JR (1980) A constructivist view of emotion. In: Kellerman H, Plutchik, R. (eds.) Emotion: theory, research and experience, Vol. 1. NY, pp 305-337 Academic, New York, NY
- Ben Ze'ev A (2000) 'I only have eyes for you': the partiality of positive emotions. *J, Theory Soc Behav*, 30(3): 341-351
- Campbell N (2003) Databases of expressive speech. In: Proceedings of Oriental COCODSA Workshop, Singapore, October 2003
- Cosmides L, Tooby J (2000) Evolutionary psychology and the emotions. In: Lewis M, Haviland-Jones J.M. (eds). Handbook of Emotions, 2nd edn. Guilford, New York, NY
- Cowie R (2005) What are people doing when they assign everyday emotion terms? *Psychol Inq*. 16(1): 11-18
- Cowie R, Cornelius R (2003) Describing the emotional states that are expressed in speech. *Speech Commun* 40: 5-32
- Douglas-Cowie E, Campbell N, Cowie R, Roach P (2003) Emotional speech: Towards a new generation of databases. *Speech Commun* 40: 33-60
- Douglas-Cowie E, Devillers L, Martin J, Cowie R, Savvidou S, Abrilian S, Cox C (2005) Multimodal databases of everyday emotion: facing up to complexity. In: Proceedings of Interspeech 2005, Lisbon, Portugal, 2005, pp 813-816
- Eysenck MW (1997) Anxiety and cognition: a unified theory. Psychology Press, Hove, East Sussex, UK
- Eysenck MW, Derakshan N, Santos R, Calvo MG (2007) Anxiety and cognitive performance: attentional control theory. *Emotion*, 7(2): 336-353
- Frijda NH (1987) The emotions. Studies in emotion and social interaction (Ser.). Cambridge University Press, Cambridge
- Gendolla GHE, Krusken J (2002) Mood, task demand, and effort-related cardiovascular response. *Cog, Emo*, 16: 577-603
- Goldie P (2000) The emotions: a philosophical exploration. Clarendon Press, Oxford
- Harré R (1986) The Social construction of emotions. Blackwell, Oxford
- Hilgard ER (1980) The trilogy of mind: cognition, affection and conation. *J Hist Behav Sci* 16: 107-117
- Huang MP, Alessi NE (1999) Presence as an emotional experience. In: Westwood JD, Hoffman HM, Robb RA, Stredney D, (eds.) Medicine meets virtual reality: the convergence of physical and informational technologies options for a new era in healthcare. IOS Press, Amsterdam, pp 148-153
- Isen AM (1998) On the relationship between affect and creative problem solving. In: Russ S. (ed.) Affect, creative experience, and psychological adjustment. Braun-Brumfield, Ann Arbor, MI, pp 3-17
- James W (1984) What is an emotion? *Mind*, 9: 188-205
- James W, Perry RB (1920) Collected essays and reviews. Longmans, Green, New York, NY
- Laney C, Campbell HV, Heuer F, Reisberg D (2004) Memory for thematically arousing events. *Mem, Cogn* 32: 1149-1159

- Lazarus R (1999) The cognition-emotion debate: A bit of history. In: Dalgleish T, Power M. (eds.) *Handbook of cognition and emotion*, Wiley, Hoboken, NJ, pp 3–19
- LeDoux JE (2000) Emotion circuits in the brain. *Ann. Rev. Psychol* 23: 155–184
- Norman DA, Ortony A, Russell DM (2003) Affect and machine design: lessons for the development of autonomous machines. *IBM Syst J* 42: 38–44
- O'Connor GF, Arnold JD (1973) *Intonation in colloquial English* 2nd edn. Longman, London
- Ortony A (2002) On making believable emotional agents believable. In: Trappl R, Petta P, Payr S. (eds.) *Emotions in humans and artifacts*. MIT Press, Cambridge, MA, p 189
- Ortony A, Turner TJ (1990) What's basic about basic emotions? *Psychol Rev*, 97: 315–331
- Ortony A, Clore GL, Collins A (1988) *The cognitive structure of emotions*. Cambridge University Press, Cambridge
- Osgood CE (1957) *The measurement of meaning*. University of Illinois Press, Urbana, IL
- Panksepp J (2003) At the interface of the affective, behavioral, and cognitive neurosciences: decoding the emotional feelings of the brain. *Brain Cogn*. 52: 4–14
- Parkinson B (1995) *Ideas and realities of emotion*. Routledge, London
- Plato (2003) *The Republic* (trans: Lee HDP), 2nd edn. Penguin Books, London
- Plutchik R (1980) *Emotion: a psychoevolutionary synthesis*. Harper & Row, New York, NY
- Rapaport D (1953) On the psycho-analytic theory of affects. *Intl. J Psycho-Anal* 34: 177–198
- Roesch EB, Fontaine JRJ, Scherer KR (2006) The world of emotions is two-dimensional.. or is it? In: HUMAINE Summer School, Genoa, Italy URL <http://emotion-research.net/ws/summerschool3/RoeschFontaineScherer-06-Genova.pdf>. Accessed on 11 October 2010
- Rolls ET (1999) *The brain and emotion*. Oxford University Press, Oxford
- Russell JA (2003) Core affect and the psychological construction of emotion. *Psychol Rev*, 110(1): 145–172
- Russell JA, Fehr B (1994) Fuzzy concepts in a fuzzy hierarchy: Varieties of anger. *J, Pers, Soc, Psychol*, 67: 186–205
- Russell JA, Feldman Barrett L (1999) Core affect, prototypical emotional episodes, and other things called emotion: dissecting the elephant. *J Pers Soc Psychol*, 76(5): 805–819
- Sabini J, Silver M (2005) Why emotion names and experiences don't neatly pair. *Psychol Inq*. 16: 11–48
- Sander D, Grandjean D, Scherer KR (2005) A systems approach to appraisal mechanisms in emotion. *Neural Netw* 18(4): 317–352
- Scherer KR (2003) Vocal communication of emotion: a review of research paradigms. *Speech Commun* 40: 227–256
- Scherer KR, Roesch EB, Bänziger T (2004) Preliminary plans for exemplars: Theory humaine deliverable D3c. Technical report, University of Geneva, 2004. URL <http://emotion-research.net/projects/humaine/deliverables/D3c.pdf>. Accessed on 11 October 2010
- Stemmler G, Heldmann M, Pauls C, Scherer T (2001) Constraints for emotion specificity in fear and anger: the context counts. *Psychophysiology*, 69: 275–291
- Sternberg RJ (1988) Triangulating love. In: Barnes ML, (ed.) *The psychology of love*. Yale University Press, New Haven, CT
- Stocker M, Hegeman E (1996) *Valuing emotions*. Cambridge University Press, Cambridge
- Tomkins SS (1964) *Affect, imagery, consciousness vol. I*. Springer, New York, NY
- Tomkins SS (1991) *Affect, imagery, consciousness vol. III. The negative affects: anger and fear*. Springer, New York, NY
- Wichmann A (2002) Attitudinal intonation and the inferential process. In: *Proceedings of Speech Prosody 2002 Aix-en-Provence, France, April 11–13, 2002*. http://www.isca-speech.org/archive/sp2002/sp02_011.pdf. Accessed on 10 November 2010
- Wilhelm P, Schoebi D, Perrez M (2004) Frequency estimates of emotions in everyday life from a diary method's perspective: a comment on scherer et al.'s survey-study "emotions in everyday life". *Soc Sci Inf* 43(4): 647–665