

humaine

D3i

Emotional life

Terminological and conceptual clarifications



Date: 1st October 2007

IST project contract no.	507422
Project title	HUMAINE Human-Machine Interaction Network on Emotions
Contractual date of delivery	<i>31st September 2007</i>
Actual date of delivery	<i>7th October 2006</i>
Deliverable number	D3i
Deliverable title	Emotional life: Terminological and conceptual clarifications
Type	Report
Number of pages	50
WP contributing to the deliverable	WP3
Responsible for task	R Cowie
Author(s)	R Cowie
EC Project Officer	Philippe Gelin

Address of person responsible: r.cowie@qub.ac.uk

CONTENTS

Overview4

Describing emotional life.....5

Towards a consensus50

Overview

This document is for people who are working in the area that is sometimes called affective computing, sometimes called emotion-oriented computing, sometimes other things. It is clear from experience in HUMAINE that they have problems with terminology. The aim of the document is to help them to deal with the problems.

The problems are not self-inflicted. They come from the fact that people in the area have to navigate at least three systems for representing the phenomena they want to talk about.

First, they have to deal with everyday language about emotion and emotion-related matters, for two reasons. They have important interactions with people who take it as read that they will communicate in everyday language – people in funding bodies, corporate user interested in manufacturing or installing systems, individual users reading manuals or asking for advice, opinion formers, etc. In addition, their systems will have interactions with people who take as read that the systems will communicate with them in everyday language (or a subset of it).

Second, they have to deal with the language developed by scientists researching emotion, affect, and related topics. Their language is derived from everyday language in the long run, but changed in some very interesting ways. To learn from the scientific community, people in emotion-oriented computing have to understand its language. To add to the problem, the scientific community is far from unified. Different groups use the same words in markedly different ways.

Third, they have to deal with an underlying system of representation that provides the medium in which people think about emotion and related matters. Intuitively, it is difficult to doubt that such a system exists, interacts richly with explicit linguistic representations, and complements them (things that we think about emotion are often very difficult to put into words). However, there is very little satisfying analysis of the system.

Because the problems are genuinely difficult, HUMAINE has approached them in a gradual way. This deliverable reflects a series of presentations to HUMAINE meetings, refined in the light of comments from both HUMAINE members and others (particular thanks are due to Brian Parkinson, Peter Goldie and Tanja Baenzinger for comments on earlier versions).

The first part of the deliverable is based on a paper which was presented to the ‘Crosscurrents’ workshop in Athens, and revised in the light of responses there.

The second part is based on a series of interactions following the Crosscurrents meeting, culminating in an interdisciplinary meeting at Haifa in 2007. Discussions there led to an agreed statement on some of the central issues raised in Athens, the ‘Declaration of Haifa’. It is rooted in a version of the ‘Emotional Life’ paper, but it extracts ideas on which there appears to be a reasonable level of consensus.

The process of reaching consensus does not end with this deliverable. Work is in progress on a chapter for the HUMAINE handbook, which will draw on some of the detail in the ‘Emotional Life’ paper, taking account of the consensus reflected in the ‘Declaration of Haifa’. It is only through sustained effort to achieve consensus and completeness that the area will overcome its conceptual and terminological problems.

Describing emotional life

This position paper is part of the theory workpackage in HUMAINE, the EC's network of excellence on emotion-oriented computing. It responds to one of the basic needs of people working on emotion-oriented computing, which is to have appropriate ways of describing the human phenomena that the area deals with.

The paper is particularly concerned with allowing emotion-oriented computing to talk about the whole range of phenomena that make up emotional life. That is a different aim from schemes that try to cover a sub-area of emotional life very coherently or precisely. They are 'depth first'; the approach here is 'breadth first'. The two are complementary rather than competitive. For instance, breadth first approaches are needed to assess how much of the relevant territory a particular depth first approach has succeeded in covering; and translating bits of a broad description into depth reveals where it has glossed over fundamental difficulties.

Part of the strategy for achieving breath is to consider the domain at different levels. The paper considers five:

- terms that cover (or might be thought to cover) the domain as a whole
- the elements of emotional life
- the states that make up emotional life
- the dimensions of emotional life
- the processes that make up emotional life.

The effect of focusing on any particular level is to bring some issues to the fore, and make it difficult to focus satisfyingly on others. Considering all five is the best practical way of minimising the range of issues that are left completely unacknowledged. That said, some levels seem to be specially restrictive. In particular, thinking in terms of states foregrounds what may well be an artificial problem – reconciling the phenomena with a type of conceptual framework that seems not to fit them more than loosely. That is a significant issue when thinking in terms of states is so likely to be the first option for so many people.

The exercise aims to develop working descriptive tools rather than a coherent, theoretically motivated perspective on emotion. At this stage, pressing for coherence is at odds with covering the whole domain. It is still the case that various ways of talking about emotion – academic and everyday – seem to capture some aspects of the domain particularly well, even if they overlap with other expressions, or carry questionable implications. Turns of phrase like that are brought into the framework if they capture something that is difficult to express in any other way. Replacing them with terms that fit together in a seamless whole is beyond the scope of this exercise.

Contrasting styles of definition are part of the diversity. In particular, writers use two extreme styles, which will be called experiential and conceptual. An experiential definition is couched in terms of properties, processes, and so on that can be observed by ordinary means (experiencing one's own state or perceiving other people's), and which is therefore relevant to clarifying the way a term is used in everyday life. A conceptual definition is couched in terms structures and systems that are supposed lie behind the observable patterns, and which can be observed, if at all, only by using specialised techniques (technological, inferential, or both). It is usually tied to a theory, philosophical, psychological, or biological.

Research in the various disciplines that deal with emotion has developed a wide range of definitions that are wholly or partly conceptual. These are intellectually interesting, and can be directly useful to some parts of emotion-oriented computing (particularly those concerned with synthesis, which can implement hypothetical mechanisms and use them to control agents). It is not only experts who use conceptual definitions – people in general are very ready to adopt ‘naïve psychological’ theories of emotion which hinge on reconceptualising it (e.g. as a visceral phenomenon in the style of James, or an inner glow in the style of Cannon). Nevertheless, in our current understanding of emotion, experiential definitions have a claim to be both practically and logically prior.

Practically, experiential definitions are what people in emotion-oriented computing need to map the terms they use onto phenomena that they might be interested in recognising or modelling, and that the public might expect them to work on. Logically, in this area, evaluating a conceptual definition involves working out whether it addresses the full range of issues that an experiential definition indicates it should. That may change as the science matures. Conceptual descriptions can become the primary way of defining concepts (for instance, the experiential definition of a metal as a shiny, ductile, and malleable solid has been displaced by a conceptual definition in terms of the electrons in the outermost shell). But theories of emotion are not yet at the level of theories in physical chemistry – as witness the many different conceptual definitions that different theorists propose for relevant terms (not least ‘emotion’ and ‘affect’).

In that situation, the paper is particularly concerned to ensure that terms have experiential definitions if possible. They are the bedrock of communication, and establishing some kind of agreement about them is essential. Without them, there is a very real sense in which we simply do not know what the terms mean, however elegant the conceptual images that they evoke may be.

A final point at this stage is that the exercise would not be necessary if everyday language provided adequate ways of describing emotional life. In reality, everyday emotion language, and the conceptual frameworks that underpin it, are about providing people with manageably simple ways of thinking and talking about massively complex phenomena. The simplification comes at a price, and the result is that everyday understanding of emotion is shot through with oversimplifications, misdirections, and traps.

Avoiding those problems also comes at a price. Readers will probably feel at various points that things surely do not need to be so elaborate and longwinded, and there is surely no need to invent all these terms (starting with ‘emotional life’). Unfortunately, there is a need. No doubt most of what this paper does could be done better. But it seems very unlikely that it can be done much more simply. Describing something as complex as emotional life simply does take a deal of effort.

1 What is to be described?

The answer proposed here is: emotional life. That means the sum total of states, processes, feelings, forms of thought, and forms of action that distinguish human life as we know it from a life we would agree was totally unemotional.

Two different answers need to be considered, which are:

- emotion in a particular, strong sense
- affect.

Everyday language makes it easy to confuse the different answers, and a sound descriptive system should prevent that. The next few sections propose a way of avoiding the confusion.

1.1 Emotion in the strong sense

It is a fact of everyday language that the word ‘emotion’ has a strong sense – it is roughly what people tend to describe when they are asked to say what the word emotion *really* means.

There is a standard kind of experiential definition for emotion in that sense, along the following lines:

a state where the person’s whole system is caught up in the way they react to a particular person or situation – which may be in reality or in their mind. Specifically,

- it involves distinctive positive or negative feelings about the people or situations involved
- it involves impulses to act or express yourself in particular ways and avoid others
- it involves distinctive changes in your body, for instance in your heart rate or tendency to sweat
- it doesn’t usually last very long – it comes on quite quickly, and dies down reasonably soon (unless there is something very unusual happening)

Conceptual definitions are more varied, but one due to Scherer gives an interpretation of that kind of phenomenon which has wide (though not universal) support. Scherer defines emotions as follows:

episodes of massive, synchronized recruitment of mental and somatic resources allowing to adapt or cope with a stimulus event subjectively appraised as being highly pertinent to the needs, goals and values of the individuals” (Scherer et al 2004, p 10)

There is no doubt that definitions like these are important for emotion-oriented computing. However, states of the kind that they describe make up a small part of emotional life. We obtained a rough estimate of the frequencies involved from unedited recordings of people taking part in a reality TV program, consisting of 9 tapes, of about half an hour each, that tracked people through a series of experiences – some demanding, some routine. A rater watched the tapes, and signalled instant by instant which of three types of state the person being recorded was in: emotion in the strong sense: unemotional; or intermediate between the two. The graph below shows the results.

The exact ratios are not important here. What matters is that the study makes a basic point. Very little of life is truly unemotional, and not much more consists of emotion in the strong sense. Most of it lies between the two.




Figure 1: the time spent by a person recorded on video in three categories unemotional; emotion in a strong sense; and elements of emotion, but not emotion in a strong sense. The top panel shows data for the whole set of recordings. The lower panels show the extremes – the tape where the target was least emotional in the left hand panel, and the tape where the target was most emotional in the right hand panel.

1.2 A three-part division

The quirks of everyday language make it easy to slip into thinking that if people are not unemotional, they must be emotional in the sense of Scherer’s definition. As Figure 1 shows, that is a recipe for ignoring most of emotional life. A good descriptive system should prevent that. Hence the paper proposes terms designed to mark the distinction that everyday language tends to obscure, as follows.

- *Emergent emotion* consists of episodes that fit Scherer’s definition and the corresponding experiential description
- *Pervasive emotion* is what runs through (or colours) the rest of emotional life
- *Emotionless* states involve neither emergent nor pervasive emotion.

Earlier versions of this scheme used the phrase ‘episodic emotion’ to contrast with ‘pervasive emotion’ (following Stocker and Hegeman, 1992). The problem with that terminology is that ‘emotional episode’ is used by several writers to describe a natural unit of emotional life, which

may or may not involve emergent emotion. The phrase ‘emotional episode’ is too useful to give up, and so, reluctantly, ‘episodic emotion’ has to be sacrificed.

It is worth noting how extraordinarily difficult it is to find suitable ways of conveying these distinctions in everyday vocabulary. In the process of looking for suitable alternatives to ‘emergent’ or ‘episodic’, all of the following terms have been considered: breakthrough, climactic, cohesive, compelling, concentrated, connected, consummated, converged, co-ordinated, crystallised, dense, dominant, flaring, focussed, foreground, fully-fledged, importunate, integrated, obtrusive, suprathreshold, unassimilated. Most of them capture some of the key distinctions, but they are bound up with connotations that tend to misdirect people. Experience suggests that people are prone to be swayed by these misdirections in spite of instructions to focus on the definitions associated with the term rather than the term itself. That is interesting, and it is taken up later.

A word which is not included in the list above, and which would in some ways be ideal, is ‘primary’. Unfortunately the term ‘primary’ has a long standing association with a conceptual definition based on the theory that in emotion, as in colour vision, some states are pure and others are produced by combining these pure states with each other and with more conceptual ingredients. The theory has no substantial support in modern research on emotion, but there is no realistic chance of prising the term loose from it in the near future.

1.3 Affect

Recently it has become common for research in computing and other areas to say it is concerned with affect. Part of the appeal is that the term ‘affect’ is so fluid. It does not have a sharply defined meaning in everyday discourse. Cynically, it has been said that ‘affect’ is “just a fancy way to say ‘feelings’” (Alleydog web glossary). In common use, that may not be far wide of the mark. Theorists have a different attraction to a word with that kind of plasticity: it allows them to propose conceptual definitions arising from their particular theoretical perspective without the awkwardness of having to align with a well-established experiential definition. Hence different conceptual definitions of ‘affect’ have been put forward by a long succession of major writers – such as Wundt, Freud (Rapaport, 1953), Hilgard (1980), Tomkins (1964, 1991), Panksepp (2003), and Russell’s group (Russell & Feldman Barrett 1999, Russell 2003).

Nevertheless, there is a core of meaning that it seems worth trying to reach agreement on. It rests on two deeply rooted ideas. First, Plato proposed a three-part division of mind which, after various shifts in translation, tends to be stated as ‘cognition, affect, and conation’. The division was drawn to the attention of modern psychologists by figures like Hilgard and Tomkins. According to that usage, it is a defining characteristic of affect is that it is distinct from cognition. Second, the dominant approach to emotion over the past few decades has been described as cognitive, because it emphasises the fact that states like fear, anger and so on involve distinctive ways of perceiving and evaluating significant events.

If the terms ‘emotion’ and ‘affect’ are used in line with those ideas, then they are obviously at different levels. In terms of Plato’s division, emotion is complex. It certainly involves affect. Contemporary psychology is very clear that it also involves cognition. It often involves conation too, at least according to Aristotle (who though Plato’s division was simplistic): he saw impulses as a key element of emotion, and impulse definitely belongs in the category of conation.

It therefore makes sense to define affect as the hypothetical element of emotion which is separate from the thoughts, percepts and impulses that are part of emotion. That suggests it might be

identified either with feelings, or with physiological systems at a lower level than those that underpin cognition, or with both. Panksepp and Russell both offer conceptual definitions that have broadly that sense, though they differ in other specific ways.

There is no doubt that a concept of that kind has intuitive appeal. However, it is definitely not a substitute for ‘emotion’ and related terms, either in terms of its meaning or in terms of ability to designate a set of phenomena that researchers know they should address before they start.

If affect is understood in that way, as the non-cognitive (and non-conative) component of emotion, then one might assume that the term ‘affective computing’ would refer to technology concerned with tapping specifically into that component (for instance using physiological measures). That seems a fair summary of what is done in the core research effort that describes itself in that way. It is an empirical question whether such an approach will be effective at dealing with the complex phenomena that clearly do play a large part in everyday life, phenomena which ‘emotion’ and cognate terms are the only convenient way of describing.

1.4 Conclusion

The intuitive case for developing emotion-oriented computing rests on the sense that there are widespread issues for it to deal with. If so, what it is concerned with is emotional life. That includes both pervasive and emergent emotion. Studying emotional life should not be equated with studying affect, assuming that affect means essentially the (hypothetical) non-cognitive element of emotion. Emotion is a complex phenomenon, and cognition is as integral to it as affect.

It is because the word emotion has that strong implication of a phenomenon with multiple elements that it does not make sense to reserve it for emergent emotion, as some authors want to do. Doing that takes away the only natural way of referring to the less tightly integrated, but still complex phenomena that make up most of emotional life.

Note that the descriptive issues here are tricky, not because people involved in research use language carelessly, but because the everyday descriptive system that they start with has problems built into it. As a result, any satisfactory set of descriptive conventions must be at odds with everyday usage in some way. However, if it is too far at odds, the result is miscommunication with the world at large – including funding bodies, users, opinion formers and lawmakers. Some kinds of miscommunication are simply awkward: others have to be seen as downright unethical, particularly those which give an inflated impression of experts’ competence because they use a word in a narrower sense than the one outsiders automatically attach to it.

Similar issues run through all the levels that will be considered, though they are only made explicit here. Bringing them into the open may increase the chances that they will be dealt with constructively.

2 Elements of emotional life

It is basic to the concept of emotion that it involves diverse kinds of element – not only affective, but also cognitive and conative; and not only within the person experiencing an emotion, but also in his or her relationships to other people and things. A useful first step in the descriptive task is to list the main kinds element that go to make up emotional life.

The list here is meant to identify issues that an adequate description of emotional life cannot ignore. It certainly not complete, and items are included or excluded on the basis of intuition that they are essential rather than well-formulated criteria. The list also reflects the concepts that arise later in the paper, and that it seems natural to introduce here. It would be useful to find ways of systematising the list and agreeing what are natural limits of this kind of description.

2.1 *Feeling*

Emotion is sometimes identified with a particular kind of feeling, but it seems fair to say that in the usual, core meanings of the terms, feeling is an element of emotion. Most people assume that they know an emotional feeling when they have one (at least when their attention is directed towards it), however specifying the exact kind of feeling that is associated with emotion is notoriously difficult. Feeling the roughness of a surface, for instance, is not intrinsically emotional. Feeling that the heart is racing and the palms are clammy is a debated case. James (1884) believed that that the feeling element of emotion was nothing but that kind of visceral feeling, but the weight of opinion since suggests that it is at most one component of emotional feeling. Elements that are listed later may be relevant to the distinctive quality of emotional feeling.

A logical alternative is that is feeling is not a component of emotion, but a mark of the cloth from which the main components of emotion are cut. On that account, the characteristic processes involved in emotion have similar functions to conscious judgment and intention to act; but the mechanisms are of a different kind. The link to feeling is that these mechanisms are, so to speak, on the same psychological level as visceral feelings or pain. On that account, emotion involves a distinctive kind of valuation which is felt rather than calculated ('felt valuation'); a kind of inclination to act which is felt rather than deliberately decided ('felt action tendency'), and so on.

2.2 *State awareness*

Being in an emotional state often includes knowing that we are in a distinctive state. It would be controversial (to say the least) to claim that it always includes knowing that we are in a distinctive state. One of the points that William James made with his famous example of the bear in the woods is that other components of emotionality may precede that kind of awareness of our state – we run first, and only later recognize that we are afraid. On the other hand, it does not seem controversial to say that emotion invites state awareness. In that it contrasts with, for instance, prejudice or rashness, whose power lies partly in the fact that it is so easy not to be aware of them.

2.3 *Partial perception*

People say that a red mist comes over an angry person's eyes, or that a lover has eyes only for a beloved. The term 'partial perception', due to Ben Ze'ev (2000), neatly captures that distinctively emotional style of perception. It has the dual implication that perception of that kind provides incomplete information and that it incorporates a positive or negative view of key things or events.

2.4 Evaluation / Valence

Emotion is deeply linked to the sense of ‘weal or woe’, as Arnold put it (1960). Emotional evaluation is bound up with partial perception, and not to be confused with rational risk assessment or quality assurance. Relevant factual information does not automatically change it, and it tends to be bound up with disposition to act – in the simplest case, to approach or avoid. It is also linked to feeling, though it can exist without being felt, as in a long term attitude to a person or thing which does not surface unless that person or thing comes to mind.

Valence is the outcome of an emotional evaluation – a sense of how much ‘weal or woe’ is associated with an object or event.

2.5 Appraisal

Theorists since Arnold (1960) have used the term ‘appraisal’ to describe a kind of structured perceptual evaluation that is characteristically emotional. The term is included here for completeness, but it is treated at length later – first in the section on states, then more fully in the section on process.

2.6 Arousal

The term ‘arousal’ refers to a generalized departure from default levels of activity, disposition to act, or preparedness to act (for instance through raised heart rate). Clearly departures from default levels of arousal are often part of emotion.

It is not self-evident that arousal is a single variable. One of the pioneering papers in the field (Schlosberg 1954) argued that different arousal-related subsystems vary independently in different emotions. The idea was sidelined by an upsurge of interest in what seemed to be a single physiological activating system. More recently, physiology has recognised that there are indeed several distinct brain systems associated with arousal.

2.7 Core affect

Russell and his collaborators use the term ‘core affect’ to identify a specific variant of the affect concept. As Feldman-Barrett (2006) puts it:

core affect is a neurophysiological barometer of the individual’s relation to an environment at a given point in time. Core affect is the ongoing, ever-changing state that is available to be categorized during emotion conceptualization. ... the neurophysiological state... is available to consciousness and is experienced as feeling good or bad (valence) and to a lesser extent as activated or deactivated (arousal).

2.8 Feeling towards

Goldie (2001) goes a step further than Feldman-Barrett. He proposes that the feelings which are characteristic of emotion are about something. That means they are neither like cognitions, which are about something but do not have the quality of feeling; nor like other types of feeling, like a headache, which are about nothing but themselves. He calls these feelings ‘feelings towards’.

2.9 Action tendency

Frijda (1986) argued that tendencies to act in a particular (biologically significant) way were integral to emotion, and that their identity was bound up with these ‘action tendencies’ – to approach (desire), to avoid (fear), to reject (disgust), and so on.

2.10 Impulsiveness / feeling like

Aristotle cited impulses to achieve revenge as a defining feature of anger. The concept of impulsiveness is linked to action tendency, but it has more specific implications which are worth keeping in mind – particularly implications about the ways action tendencies interact with (and interrupt or subvert) deliberate thought.

Everyday language has a phrase that captures this element in an interesting way. It is ‘feeling like’, as in ‘I feel like screaming’, or ‘I feel like dancing’. What is interesting is that the phrase suggests the action tendencies or impulses are part and parcel of the domain of feeling. In other words, feeling extends not only to objects (as in ‘feeling towards’) but also to actions (as in ‘feeling like’).

Emotional activity

This is a blanket term to cover the many voluntary actions and involuntary changes that an observer may register as signs of emotion. The next three headings reflect the main types.

2.11 Involuntary changes

This term covers phenomena like blushing, sweat, trembling, gasping, and other movements or changes that are not the result of a conscious intention to act in a particular way.

2.12 Emotional gestures

This term is used to cover actions that do not generally occur unless some specific circumstances exist, one of which is the presence of an emotion. Facial gestures such as smiles and frowns are obvious examples. Some interjections, hand gestures and body postures also belong in the category.

2.13 Emotional modes of action

This term is used to cover signs of emotion that reside in the way actions are executed rather than the fact that they are executed at all. Emotional ways of speaking or standing are salient examples. However, the category is much wider. For instance, just as impulses are a feature of emotional thought, so impulsive behaviour is a feature of emotional action.

2.14 Modified cognition

Emotion impinges on various aspects of information processing, including logical reasoning and use of language. It is tempting to use the word irrationality to describe that element, but that has negative connotations which are unhelpful. It is not irrational in the ordinary sense to interrupt contemplation when William James’ famous bear appears yards away. The phrase ‘hot cognition’

has gained currency, but it fits some kinds of emotion much better than others. ‘Modified cognition’ is clumsy, but not misleading.

2.15 Priority

The sense that events or things are important or unimportant is related to evaluation (if something is very bad, dealing with it probably should be a priority); and it is also central to emotion. As with evaluation, the emotional sense of importance is distinct from rational prioritisation.

2.16 Engagement

It seems not to be discussed a great deal in the recent literature, but intuitively it seems natural to contrast emotion with detachment. The sense of being personally engaged with, and vulnerable to, events around, seems integral to emotion, and characteristic of it. Russell (1976) has made a related point under the heading ‘depth of experience’, and research on presence tends to assume that emotion is somehow bound up with feeling that an experience is ‘for real’. Ortony (2002) used the term ‘caring’ to express what seems to be a related concept. It is also linked to the concept of ‘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 & Alessi 1999).

2.17 Intentionality

Mental states are described as intentional if they are directed towards or concerned with an object (as against being self-contained). Not all emotion-related states are intentional, but it is integral to the character of some states that they entail a particular kind of intentionality.

2.18 Object

The object of an emotion is the person, thing, situation or whatever with which the emotional feelings are concerned – the thing that is frightening in the case of fear, the person who is beloved in the case of love.

2.19 Trigger

The trigger of an emotion is the person, thing, situation or whatever that precipitates an emotional event. It may be the same as the object of the emotion (e.g. the famous bear both precipitated James’ fear and was its object), but it need not be (e.g. finding an old letter may precipitate an episode of grief about the person who wrote it).

2.20 Recipient

The person to whom emotional expressions are directed. Particularly with interpersonal emotions, specifying the recipient is part and parcel of the specifying the emotion. For instance, expressing contempt to person X does not imply contempt for the others who are present.

2.21 Audience

People who are known to be present when an emotion is expressed, but who are not recipients. There may be a sense in which an audience is always a kind of recipient – signs are meant to give

them a message, even if it is the same message as the recipient. But intuitively it seems there is a distinction worth making, even if it may be subtler than one might assume at first sight.

2.22 Display rule

This is a concept due to Ekman, which postulates that external expression of an emotion is a function of a 'true' internal state and socially determined 'display rules' which specify how it may be expressed in the context of a particular audience. It has been criticised for artificially separating a pure (but hidden) emotion from social factors governing its expression. Of course there are sometimes sharp dissociations, but it usually seems as if the emotions that we experience in social contexts are much more fundamentally shaped by the others involved and our relationships to them.

2.23 Emotion label

An everyday word that is used to categorise an emotional or emotion-related state.

2.24 Emotion percept

The impression that one person forms of another's emotional state and orientation. It may include an emotion label, but it is rarely exhausted by one.

2.25 Empathy

Feeling in one person which he/she believes mirrors the emotions being experienced by another. It certainly may be part of an emotion percept: arguably it always is.

3 Emotional and 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 a word like emotion in the strong sense 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 certainly provides a useful framework.

3.1 The order of personal states

People are quite reasonably sceptical when emotion and related terms are used as a catchall for everything that is conveyed in non-verbal communication. It is important to be clear that states which can reasonably be thought of as emotion-related are only part of a larger domain. A natural name for it is ‘personal states’. Included at that level might be states of health (pains, tremors, infections, etc) and wellbeing (vigorous, exhausted, etc); social states (involving affiliation or relative status); and cognitive states (confusion, preoccupation, etc).

A very large proportion of personal states have at least emotional connotations. For example, courage seems to be essentially a behaviourally defined state; and yet it is natural to assume that courage will usually have intimate links to emotional factors (such as the way fear is suppressed or over-ridden). Should courage, then, be considered an emotion-related state?

There is probably no absolute way to answer that kind of question. Pragmatically, perhaps a state should be considered emotion-related if dealing with it in a satisfactory way depends on modelling the emotional elements. 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.

3.2 The family of emotion-related states

On the taxonomic metaphor, emotion-related states make up a family, and can be divided into genera. Some genera are very well-established, such as moods. There are other terms it is natural to think of as being roughly at the same level, but it is not self-evident how they all fit together.

3.2.1 Scherer’s generic matrix

At the generic level, a matrix due to Scherer provides a well-defined starting point. Figure 2 shows an adapted version, modified partly to keep terminology consistent, and partly to bring out an important issue, which is the centrality of emergent emotion. What characterises the other genera is that they have some, but not all of the features associated with emergent emotion.

The design features that Scherer lists are partly conceptual and partly experiential. Synchronisation is central to the conceptual definition of emergent emotion that was given earlier. Appraisal is also conceptually defined, and particularly associated with emergent emotion.

‘Event focus’ is concerned with the object of the emotion. The contrasts involving it have an experiential basis, but tend to be developed in a conceptual direction. In particular, mood is routinely defined as affect that has no object. It is certainly an elegant distinction, but it is less obvious that it actually fits experiences that people usually call moods.

'Design features'	Emergent emotion	Interpersonal Stances	Moods	Attitudes	Affective disposition
Impact on behaviour	++	+		+	+
Intensity	++	+	+	+	
Rapidity of change	++	++	+		
Brevity	++	+	+		
Event focus	++	+			
Appraisal elicitation	++				
Synchronisation	++				

Figure 2: genera of emotional state (after Scherer 2004).

Features in a third group are to do with intensity and timing. Emergent emotions are intense and transient (short, with rapid onset and offset). The same features make up the core of the qualities listed for moods and stances, but to a lesser degree. They might be called envelope features, since they define a temporal envelope, but not what goes in it.

The remaining feature, impact on behaviour, relates to the concept of action tendency. Apart from mood, the other states are associated with some kind of action tendency, but to a lesser extent than fullblown emotion.

A later table proposes strictly experiential definitions of the states, along with others.

Scherer's generic matrix is a model of the kind of order that a taxonomy of emotion-related states might aim for. However, if one is concerned with completeness, it is a starting point rather than a finished product. There are states that it is natural to see as part of emotional life that do not have a comfortable home within the framework. The next section sets out examples.

Although the table does not provide a complete taxonomy, there is a different kind of descriptive task that perhaps it does address completely. A large body of work indicates that membership of a psychological category tends to behave as if it depended on distance from a central point in semantic space – a 'prototype' (Rosch 1978). There is direct evidence that some emotion-related concepts fit that model (Russell 1991).

A case could certainly be made that Scherer's table captures the prototypes round which people categorise emotion-related states at a generic level. The evidence described in the next section is the kind of raw material that might be used to develop a case like that. However, it is important not to confuse the proposal that these are critical landmarks in the domain – which they may well be – with the idea that very much of emotional life lies very close to these landmarks. The next section describes an approach to collecting evidence on the issue.

A final point about prototypes reflects an observation that was made earlier. It is natural to see emergent emotion as a higher-level prototype – the various states listed are considered to be emotion-related insofar as they have significant features in common with it.

3.2.2 O’Neil’s study

There are two main kinds of reason to think that Scherer’s table is incomplete if it is considered as a taxonomy. From a conceptual stance, the earlier list of elements suggests features that might be as relevant to distinctions at this level as those that Scherer lists as design features. From an experiential stance, both work with naturalistic records of emotion (Douglas-Cowie et al 2005) and discussions based in other traditions (particularly philosophical) highlight states that do not fit the categories in the table comfortably.

Those two lines of approach led to an expanded version of Scherer’s list of generic types. Early versions were presented both to people involved in emotion research and to members of the public. On that basis, they were refined by removing categories that people found difficult to understand or to distinguish from others, and by adding categories to cover experiences that informants felt were part of emotional life, and that seemed not to belong to any of the available types. The eventual list was expanded in four main areas.

1. Variants of emergent emotion. The term ‘fullblown emotion’ was reserved for the archetypal example of an emergent emotion, where the emotion runs unchecked (following Scherer, 1999 and Cowie and Cornelius, 2003). ‘Suppressed emotion’ was used for the case where the inner state is characterised by the kind of synchrony that Scherer describes, but its outward expression is deliberately controlled. ‘Simmering emotion’ was used for the case where some of the processes that Scherer describes are active, but they are not fully synchronised (and may or may not reach that point).
2. Variants of mood The term mood was reserved for states which, as in the classical definition, do not have an object. Empirically there seems to be states which are long lasting like moods, and which people may call moods, but which ‘lock on’ temporarily to one object after another. These were named ‘topic shifting’ emotions (lay users were less comfortable with an earlier name, ‘flitting’).
3. Established emotion The philosophical literature in particular is much concerned with states such as grief or shame, which are part of a person’s makeup for a long time, and come to the fore in acute episodes (often because of a relevant trigger). The term ‘established emotions’ was introduced to describe these enduring states.
4. Dimension-related states A good deal of emotion-related variation seems to involve dimensions other than valence – how aroused a person is, whether they feel in or out of control, how seriously they take things going on around them. It is not particularly easy to express these issues in a taxonomic framework, but users seemed reasonably comfortable describing them as altered states of arousal, control, and seriousness respectively.

As these comments indicate, there is an extended set of design features implicit in the list that was developed. However, it is not set out in this paper, mainly because proper discussing of the issues would be quite long.

When the list appeared to be reasonably stable, it was tested in a larger study to verify whether the categories, new and old, made sense to lay people, were common enough to be of interest, and were felt to be practically important. Users from a local community were taken through the descriptions given in the table below (Figure 3). They were deliberately not University staff or students, and the descriptions (including a supporting example) were carefully framed to be understood by them.

When participants had gone through the whole list, they were asked to go back to the categories one by one, and to answer the following questions for each one:

- How clear did you find the definition?
- How important is this kind of state to you?
- When did you last experience a state of this kind?
- Could you describe that experience briefly in your own words?

Great care was taken over ethical and data protection issues, but they are not discussed here. The free descriptions in particular are an extremely rich source, and are available for research. What makes them valuable is not simply that they are free descriptions, but that participants were directed to consider more than the relatively narrow range of types of experience that standard terms tend to evoke.

Figure 3: generic categories used in O’Neil’s study (definitions designed to be understood by naïve participants followed by examples)

<p>1) Attitudes</p>	<p>This refers to long-lasting positive or negative orientations towards things or people, which are more or less permanently part of the way you interact with them or think about them. They may just be positive or negative, or subtler.</p> <p>Example</p> <p>“I have a long lasting attitude towards nature, It never ceases to amaze me, I look at it is as being beautiful”</p>
<p>2) Established emotion</p>	<p>This refers to long-standing states that influence a person’s emotional life over weeks or years. They are linked to some person or situation that has relatively intense significance for the person. They tend to be in the background for most of the time, but they can be ‘triggered’ in a way that produces surges of overt emotion.</p> <p>Example</p> <p>“Sadness over the death of a colleague, comes through in discussion among friends remembering happier times”</p>
<p>3) Emergent emotion (full-blown)</p>	<p>This refers to states where the person’s whole system is caught up in the way they react to a particular person or situation – which may be in reality or in their mind.</p> <ul style="list-style-type: none"> – it involves distinctive positive or negative feelings about the people or situations involved – it involves impulses to act or express yourself in particular ways and avoid others – it involves distinctive changes in your body, for instance in your heart rate or tendency to sweat – it doesn’t usually last very long – it comes on quite quickly, and dies down reasonably soon (unless there is something very unusual happening) <p>Example</p>

	<p>“Fear – noise in the house at night – heart rate increased – immediately think of worst – someone in house”</p>
<p>4) Emergent emotion (suppressed)</p>	<p>This refers to states where the person’s inner feelings are like a full-blown emotion, but they are kept under control by a deliberate effort not to let them take over or show to other people.</p> <p>Example</p> <p>“Picking up brother and his friend. I took a wrong turn & friend called me an idiot. Wanted to shout back but knew it would be too much trouble so decided to bite tongue instead.”</p>
<p>5) Moods</p>	<p>This refers to states where a person’s whole outlook is ‘coloured’ by a particular kind of feeling, but it is not focused on any particular thing or event – it is just how the person feels in him- or herself. These states tend to be moderately long lasting (hours or days).</p> <p>Example</p> <p>“Depressed – no real reason, but felt down, not in mood to speak to anyone, felt everything was wrong.”</p>
<p>6) Partial emotion (topic shifting)</p>	<p>This refers to states between a mood and an emergent emotion, where a person is in a single kind of emotional state, but it is not fixed on a single thing or issue – instead the emotion attaches to one object, then another, then another, in quite rapid succession.</p> <p>Example</p> <p>“Angry and tired from previous day. Getting annoyed when styling hair because it wouldn’t stay. Angry at dog trying to get out, and things at work.”</p>
<p>7) Partial emotion (simmering)</p>	<p>This refers to states where a person experiences some of the elements of an emergent emotion, but they have not come together to form the kind of emotional state that takes control of a person’s thoughts and feelings.</p> <p>Example</p> <p>“A steward on our plane found a passenger using a mobile phone. She kept arguing and he was polite but his face coloured as if he could easily tip over the edge and get angry with her.”</p>
<p>8) Interpersonal Stances</p>	<p>This refers to relatively short-lived states where feelings towards another person come to the fore and incline you to behave in a particular way towards him or her.</p> <p>Example</p> <p>“Supportive, felt protective towards someone”</p>
<p>9) Altered state of</p>	<p>This refers to states whose main feature is that your energy levels feel</p>

arousal	<p>different from usual. It could be mental or physical energy or both. Examples are energetic, drowsy, dull, alert, withdrawn.</p> <p>Example</p> <p>“Exam time. State of alert & energy dramatically increased – regardless of how little sleep I had.”</p>
10) Altered state of control	<p>This refers to states whose main feature is that you feel either more in control of things than you usually do, or less in control. Various other kinds of feeling may be linked to the sense of control, e.g. positive or negative feelings.</p> <p>Example</p> <p>“Worried, panicky, frustrated at being let down at last minute by someone I was relying on.”</p>
11) Altered state of seriousness	<p>This refers to states where you take things more seriously than usual or less seriously than usual.</p> <p>Example</p> <p>“Tidying my room and fussing about where I ought to put an old souvenir when I suddenly realised it was totally ridiculous and had a good laugh at it all.”</p>

Figure 4 shows the main numerical results that are relevant here. Subjective ratings indicate that the definitions were generally understood. These were matched against ratings by experts who judged whether the free descriptions captured the right kind of experience, and the two showed reasonable agreement: attitude, topic shifting and simmering were the only cases where there was an appreciable problem, and even there a clear majority did understand. In terms of significance, the classical focus of research, fullblown emotion, was rated least important, with mood, attitude and established emotion most important. Frequency shows a related picture: attitude and mood had on average been experienced most recently, fullblown emotion least.

The main point to be made is that the new categories allow relatively common and important parts of emotional life to be described. Conversely, holding to familiar categories deflects attention from parts of emotional life that informants regard as relatively common and important.






Figure 4: main results of O’Neil study. Panel A shows the proportion rating clarity good or very good; Panel B shows how important each state was rated (4=’basic to the way I am’); Panel C shows time since the last experience in that category (on a logarithmic scale, 1=10 days, 2=100 days, etc).

The question arises whether adding more categories of the same kind would bring significantly more of emotional life within the scope of the description. The descriptions that O’Neil collected include examples that suggest some categories could usefully be added, but that is not for this paper. There is also a more general issue to consider, related to a point that has been made several times about the importance of prototypes.

According to the data, the average time since experiencing any of the commoner types of state is about a week. Given the number of different categories, it seems to follow that it tends to be about a day between emotion-related experiences of any kind. It is unlikely that participants had actually been completely unemotional for on average a day before the study. Presumably what happened was that they looked back to identify a state that was sufficiently different from an average baseline to say that it qualified as something different – ie one that conformed to an entrenched prototype of what counts as an ‘emotional state’. The granularity may be finer, but the study is still tapping landmarks in the emotional landscape rather than the continuous ground between them – presumably because people are so strongly drawn to think about emotion in terms of landmarks.

Two points follow. The less important is that subsequent studies should bring in different methods (like the ambulatory studies of Wilhelm & Perrez , 2004) and Wilhelm, Schoebi. & Perrez, 2004). The more important is that anyone who thinks about emotion, as a lay person or a research scientist, is likely to be drawn subtly or strongly towards a conception organised around prototypes.

3.2.3 Concluding generic taxonomy

Intuitively, it seems that emotion-related states can usefully be subdivided at a generic level. Familiar descriptions at that level are certainly useful, but there are quite large and significant parts of emotional life that they leave out. Providing a fuller taxonomy at that level is a real challenge.

The free descriptions generated by O’Neil’s study provide a valuable empirical base for extending the description. An extended set should then be tested in the same way, and so on iteratively until the set stabilises. Nevertheless, the categories that are available already do seem to cover a very substantial part of emotional life.

It is also worth noting that usages at this level seem to be extraordinarily variable. For instance, many of O’Neil’s subjects clearly have in mind a meaning of the term ‘attitude’ that is almost unrelated to the term used in a very large psychological literature, and presentations at conferences indicate that some academics have other, equally divergent meanings in mind. People involved in research should be particularly careful not to assume that their use of words at this level is either normative or the same as other people’s.

3.3 Species of emotion-related state

When people think about describing emotion, their first resort is usually the emotion labels that describe species of emotional state – fear, anger, and so on. There is a very large literature about that level of description, and this section does not attempt to review it. It has two aims. One is to set out lists of words that seem likely to be useful to emotion-oriented computing. The other is to put those lists in context, so that it is as easy as possible to see what kind of tool they are.

3.3.1 Short lists

There is an enduring hope that a short list of emotions which are ‘basic’ or ‘primary’ can span the whole domain of emotion. Cowie and Cornelius (2003) bring together some of the attempts to do that, and explain why the project does not look promising.

Much of what is attractive about the idea is captured by Scherer’s concept of modal emotions, which returns yet again to the theme of landmarks or prototypes. Scherer argues that the words that tend to be thought of as ‘primary’ or ‘basic’ emotions correspond to emotions that are relatively common because the combinations of events that evoke them are relatively common. These highly populated areas of emotion space are then treated as reference points, and serve as the centres of fuzzy classes. Emotion-oriented computing seems to have gravitated towards a similar conception in the way it has used cover classes. Labelling (for instance) is carried out using a relatively long list of terms, and they are then aggregated into a more manageable number of sets. The reduction to a short list is understood as a pragmatic device rather than depending on any very deep theory.

3.3.2 Middle sized lists

Figure 5: Cox’s composite of lists from Geneva, LIMSI and Belfast

<p>Negative & forceful</p> <p>Anger Annoyance Contempt Disgust Irritation</p> <p>Negative & not in control</p> <p>Anxiety Embarrassment Fear Helplessness Powerlessness Worry</p> <p>Negative thoughts</p> <p>Doubt Envy Frustration Guilt Shame</p> <p>Negative & passive</p> <p>Boredom Despair Disappointment Hurt Sadness</p> <p>Agitation</p> <p>Shock Stress Tension</p>	<p>Positive & lively</p> <p>Amusement Delight Elation Excitement Happiness Joy Pleasure</p> <p>Caring</p> <p>Affection Empathy Friendliness Love</p> <p>Positive thoughts</p> <p>Courage Hope Pride Satisfaction Trust</p> <p>Quiet positive</p> <p>Calm Content Relaxed Relieved Serene</p> <p>Reactive</p> <p>Interest Politeness Surprise</p>
---	---

HUMAINE partners have been active in trying to develop lists of emotional states that are short enough to be practical, but cover most of the ground emotion-oriented computing might want to consider. To see various lists that have been proposed,

go to WP6: Interaction
 then to Emotional Representation Language - Work-Folder
 then to List of emotion descriptions

Figure 5 is based on a combination of key HUMAINE sources – the Geneva emotion wheel; lists developed for the EmoTV database at LIMSI; and lists based on ratings of the importance of emotion descriptors in Belfast (reported in the 2005 special edition of *Neural Networks* on emotion). In line with earlier comments, the cover classes have no very deep motivation – they are simply a convenient division of the terms.

There is reason to think that a vocabulary of about 50 words is enough to capture a good deal of emotional life. In a series of studies using the Belfast naturalistic database, raters who watched each clip chose from 40 standard terms to describe the emotions they had seen, and were allowed to add others if they felt the standard list was not adequate. Only 10 words were added more than twice.

The databases now available in HUMAINE are more wide-ranging, and in particular they cover more of the genera described by O’Neil. It makes sense to repeat the kind of exercise that has just been described using the new databases, taking the list in figure 5 as a starting point.

3.3.3 Appraisal-based descriptions

Listing emotion-related words is one task, specifying their meanings is another. Methods based on appraisal provide the best-developed way of describing meaning systematically. Probably the best known appraisal-based analysis of emotion terms is due to Ortony, Clore and Collins. Scherer’s group have developed a richer scheme. Figure 6 shows how it analyses the meanings of 14 emotion terms

Figure 6: Analysis of emotion terms from Sander et al (2005) using Scherer’s appraisal scheme.

	Enjoyment	Elation	disgust	confident	sad	despair	worry
NOVELTY							
- Suddenness	low	hi/med	open	open	low	high	low
- Familiarity	open	open	low	open	low	v low	open
- Predictability	medium	low	low	open	Open	low	open
INTRINSIC PLEASANTNESS	high	open	v low	open	open	open	open
GOAL SIGNIFICANCE							
- Concern Relevance	open	self/relationships	body	relationships/order	open	open	body/self
- Outcome Probability	v high	v high	v high	high	v high	v high	medium

- Expectation i	consonant	open	open	open	open	dissonant	open
- Conduciveness i	conductive	v con	open	open	Obstruct	obstruct	obstruct
- Urgency	v low	low	medium	low	low	high	medium
COPING POTENTIAL							
- Cause: Agent	Open	open	open	other	open	other/ nature	other/ nature
- Cause: Motive	intent	cha/int	open	intent	chance/ negligence	chance/ negligence	open
- Control	open	open	open	high	v low	v low	open
- Power	open	open	open	low	v low	v low	low
- Adjustment	high	medium	open	high	medium	v low	Medium
COMPATIBILITY STANDARDS							
- External	open	open	open	v low	open	open	open
- Internal	open	open	open	v low	open	open	open

NOVELTY

- Suddenness	high	low	high	v low	low	open	open
- Familiarity	open	open	low	high	open	open	open
- Predictability	low	medium	low	v high	open	open	open
INTRINSIC PLEASANTNESS	low	open	open	open	open	open	open
GOAL SIGNIFICANCE							
- Concern Relevance	body	open	order	body	self	relationsh ips/ order	self
- Outcome Probability	high	v high	v high	v high	v high	v high	v high
- Expectation i	dissonant	open	dissonant	consonant	open	open	open
- Conduciveness i	obstruct	obstruct	obstruct	open	open	high	high
- Urgency	v high	medium	high	Low	high	medium	low
COPING POTENTIAL							
- Cause: Agent	other/ nature	open	other	open	self	self	self
- Cause: Motive	open	intent/ negligence	intent	open	Intent/ negligence	intent	intent
- Control	open	high	high	medium	Open	open	open
- Power	v low	medium	high	medium	open	open	open
- Adjustment	low	high	high	high	medium	medium	high

COMPATIBILITY STANDARDS							
- External	open	low	tow	open	open	v low	high
- Internal	open	low	low	open	v low	v tow	v high

3.3.4 States and words

It has been tacitly assumed so far that each emotion word describes a species of emotional state. The literature has challenged that assumption in various ways.

On one side, it is widely recognised that some words describe more than one state. The best known example is that the word anger describes two very different kinds of state, described as ‘hot’ and ‘cold’. Kahneman (1995) has argued that regret also describes two kinds of state, one ‘hot’ and one ‘wistful’. Sternberg’s manifold subdivision of love is well known (1988).

On the other side, it has been argued that different words may refer to essentially the same kind of mental state, but contextualise it differently. For instance, Sabini and Silver (2005) argued that shame and embarrassment are not names of different emotional states. The difference lies not in the emotional state, but in whether the person believes that the key event reveals a real flaw in his or her makeup. In the case of shame, he/she does; in the case of embarrassment, he/she does not.

It is easy to shrug off points like these as minor details, but they point to obstinate difficulties in dividing emotional life up. Are there really two separate angers involved if I respond immediately to an insult with hot anger, and on reflection settle into cold anger? In what sense can I imagine surgically separating the hypothesised emotional core of shame from the cognitive surround that makes it shame rather than embarrassment – and what would be the significance if it were shown to be possible?

The common element in these difficulties is that emotional states seem not to separate neatly either from others in a temporal sequence, or from connected parts of mental life. As a result, although it may turn out that models based on neatly bounded states work reasonably well, it is not something that can be assumed a priori. That links to issues discussed in the next section.

3.4 Limits of hierarchical decomposition

Taxonomic models are useful up to a point, but there are senses in which they simply do not fit the area. This section picks out one of the major problems, which is that the words listed in figures 5 and 6 are not naturally nested under the generic terms listed in figures 4 and 5.

There are species of emotional state that are strongly associated with particular genera. Feeling blue usually describes a mood, feeling friendly describes an interpersonal stance, feeling irritated usually describes a simmering emotion. However, there are also ‘species’ that cut across a great many ‘genera’. Anger can be an emergent emotion, an interpersonal stance, or an established emotion; a person can also be in an angry mood, or have an angry disposition. In addition, episodes of fullblown anger can arise out of an angry mood or an established emotion of anger.

One response is to conclude that mistakes must have been made in setting up the descriptive system (the genera are poorly chosen, the words do not truly describe emotional states); and if the mistakes

were put right, the phenomena would inevitably fall into a taxonomy. It seems more reasonable to conclude that the phenomena are genuinely fluid and interdependent in a way that means taxonomy can only ever be an approximate framework.

3.5 Describing an individual instance of emotion

There is a sense of discontinuity between the problems of describing emotion in abstract and the problems of describing specific instances of emotion. Describing a specific instance, there are overwhelmingly obvious questions – what is the object of the particular emotion? How does the person feel about it? what background is involved in making the person feel that way? and so on. The person’s circumstances and history come to the forefront, and the abstractions look less of a problem.

The sense of discontinuity probably reflects the way people are set up to understand emotion. They believe, rightly or wrongly, that given the context, they can judge intuitively how a person would feel. ‘How would you feel if ...?’ tends to be a rhetorical question because people assume, rightly or wrongly, that if the circumstances are explained, the answer will announce itself.

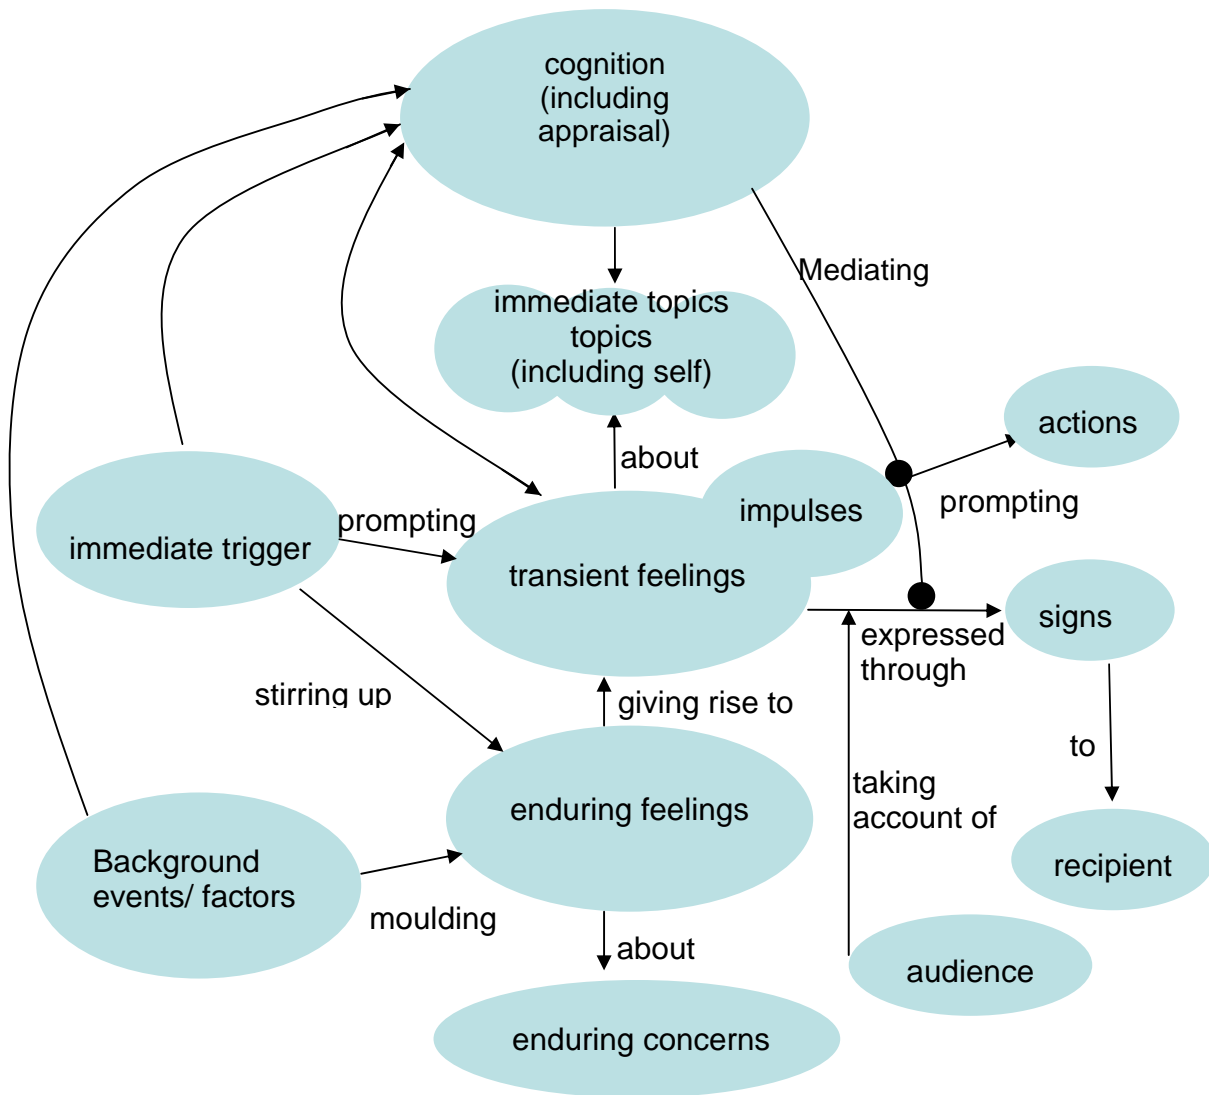
Figure 7 sketches a structure that seems to be sufficient to express most of what it is natural to say about an individual emotional episode. It arises out of work on coding samples from naturalistic emotion databases (Devillers et al 2006). It could be expressed as a form, with slots to be filled in for each component, but that tends to obscure the connections between the items.

Not all the components are always relevant. In simple reactive situations, such as James’ bear in the woods, the lower components may not be relevant; in the standard model of pure mood, only the transient feelings are relevant; and so on.

Of course, the sketch ducks central issues. In particular, it ducks the task of describing the feeling components, and their relationships to the more cognitive aspects of appraisal. At one extreme, one might think of feeling along the lines of Russell’s ‘core affect’ proposal, where feelings are essentially dimensional; at the other, one might want to think of feeling as a representation of the key items in a distinctive (and in some senses primitive) medium.

These are intriguing and theoretically central questions. However, in practice, for many parts of emotion-oriented computing, it may often not greatly matter how content is divided between relatively transparent, cognitive representations; and relatively opaque, purely emotional representations. The main issue is to avoid forcing descriptions into a framework that is too restrictive to address key issues. The deeper modelling issue is linked to discussions of high- and low-level modelling being carried out in HUMAINE’s WP7.

Figure 7: frame for describing an individual instance of emotion



3.6 Conclusion: levels and taxonomies

The model of a taxonomy provided a framework for this section. The exercise of evaluating it would be worthwhile even if only because the model has such a strong intuitive appeal, and in fact it fits well enough to be a useful approximation. However, the metaphor of a taxonomy seems to be limited in ways that it is important to recognise.

First, when people think informally about emotion, what seem to be taxonomic schemes are actually prototypical. The concepts that they use specify landmarks in the domain, and leave it to unstated rules of interpolation to cover the ground around them. Prototype-based representations may suit the area very well, but they can cause problems if they are used carelessly. In particular, it is easy to slip into the assumption that the landmarks are the proper targets of research, and as a result to end up ignoring most of the domain.

A side effect of that tendency affects the exercise of finding names for generic classes. The point was made early on (section 1.2) that it is extraordinarily difficult to find brief labels that work effectively with a short description. People are unduly swayed by connotations of the label that are not meant to be significant, apparently because they are automatically drawn to identify a very particular kind of state that they then assume is the target being considered. For instance, it would be better in some ways to describe what have been called ‘fullblown emotion’ and ‘simmering emotion’ as ‘flaring emotion’ and ‘flickering emotion’. But ‘flaring’ and ‘flickering’ have slightly threatening connotations, which bias people towards thinking in terms of negative states. It is difficult to over-ride that in the associated description, and that begins to make sense if one assumes that the intuitive frameworks people use to represent emotion hinge on specifying very particular prototypes that can then serve as the centre of a class.

Second, although it is natural to think of everyday divisions being at different levels (genus and species), it is probably not very accurate. Rather they represent different, and inter-related kinds of variation that may occur within the general kind of framework shown in figure 7. ‘Genus’ differences are mainly at a structural level, concerned with the components that are active and the relationships between them; and ‘species’ differences are mainly to do with the contents of active modules. However, certain kinds of content depend on the involvement of multiple components, and so the two kinds of variation are not independent.

The same kind of interdependence arises with another feature that has been used in a taxonomy, which is related to the ‘envelope’ design features in Scherer’s generic matrix. In one of the first serious psychological studies of emotion, Wundt (1897) divided the field into three categories based on timing - *sudden, irruptive* emotions, such as surprise, fright, and rage, which reach their highest point rapidly and then die away; *gradually arising* emotions such as anxiety, doubt, expectation, and some kinds of joy or anger; and *intermittent* emotions, that rise and fall alternatively, such as love or remorse. It is not accident that emotions have these different profiles: it is related both to the components that are active and to their contents. Correspondingly, Wundt’s organising principle is no more independent of the others than they are of each other.

4 Dimensional descriptions of emotion

Generally speaking, dimensional descriptions of emotion emerge from standard statistical analyses that have been applied to data on emotions. Raters make numerical judgments (sometimes binary) about a set of emotion-related items. The numbers are then subjected to standard algorithms (PCA, dimensional scaling, etc) which show how much of the variation in the original dataset can be explained in terms of a relatively small number of dimensions. Psychological judgments have a restricted role – in choosing the original items, and then in deciding how many extracted dimensions to take seriously, and what to call them.

One form of dimensional description is well established in emotion-oriented computing. It involves two dimensions, evaluation and activation. Essentially similar dimensions are sometimes given different names, and essentially the same space is sometimes described in terms of different axes, which are called positive and negative activation. PANAS (Positive and Negative Activation Scales) is a standard instrument for measuring affect on those axes.

Two more dimensions are very widely used to describe emotion. One is most often described as potency, though there are substantially ways of glossing essentially similar statistical patterns (e.g. approach-avoid). The other is intensity.

These dimensions have a special conceptual place, because theorists have proposed that they specify the (hypothetical) affective core of an emotion. Osgood (Osgood, Suci, & Tannenbaum 1967) made very far-reaching claims about ‘the universality of the affective E-P-A components’ (E-P-A = evaluation – potency – activation). He believed that they reflected an affective layer which was integral to semantics in general, not just to overtly emotion-related terms. SAM (Score A Mannikin) is a standard instrument for measuring affect on those dimensions. Russell and Feldman Barrett (1999) offered a narrower conception, suggesting that core affect at the centre of an emotional episode can be characterised on just two dimensions, evaluation and activation. Leaving potency out is statistically reasonable: it tends to account for much less variance than the other two dimensions, and is much less stable across studies.

Intensity is rather a Cinderella dimension, probably for methodological reasons. One of the main motives for dimensional research is to find a space that discriminates emotion words satisfyingly. Emotion words tend to be relatively noncommittal about intensity (to say ‘I was afraid’ does not specify whether the emotion was weak or strong), and so intensity does not contribute much to discriminating words (as Wundt noticed). However, emotion states clearly do vary enormously in intensity, and for emotion-oriented computing, it has a claim to be the one dimension whose importance is not seriously in doubt.

4.1 Additional dimensions

In its heyday, statistical research on emotion words threw up a great many dimensions. Proposals taken seriously in a review by Russell (1978) include dominance, aggressiveness, affiliativeness, and locus of causation. Averill (1980), writing around the same time, added depth of experience, which differentiates loving or spiritual feelings from more shallow feelings such as giddy, peevish or coy.

Recently, different concerns have thrown up different candidates. For instance, research on virtual reality has pointed to differences in quality between the emotions that arise when situations are

experienced as simulations, and the emotions that arise when the same stimuli are experienced as part of the real environment. That suggests an engagement/detachment dimension is part of emotional life, with engagement as a characteristic of compelling emotion.

Large parts of the appraisal literature also involve dimensions, though they tend not to be described in those terms. Considering the items listed in table 6, most of the features can be understood as dimensions – suddenness, familiarity, predictability, concern relevance, outcome probability, expectation, conduciveness, urgency, control, power, adjustment, and compatibility with standards (external and internal). The fundamental distinction between these and the more traditional dimensions is taken up in the next section.

4.2 What do dimensions apply to?

Arguably the most characteristic feature of accounts that are normally called ‘dimensional’ is not that they use dimensions – it is that the dimensions are used to describe a global state of the subject experiencing the emotion. That is explicit in descriptions like Barrett-Feldman’s (given earlier) which portrays affect as a barometer of the organism’s state; many more accounts are less explicit, but difficult to understand in any other way.

It has already been pointed out that many of the descriptors in appraisal theories involve dimensions, and in fact many of the dimensions in question are also found in classical dimensional analyses. A large part of the difference is that in appraisal theories, the dimensions are attached to an object. Global pleasantness is attached to an object, whereas positive affect is attributed to the person.

Coding naturalistic data gives a strong incentive to take a step further. Appraisal theories tend to take it for granted that an emotion has one, reasonably obvious object. In naturalistic data, that is often not true. In a clip from the Belfast naturalistic database, a woman talks about her pregnancy, which is a positive event. But she is painfully aware that she lost her last baby, which is negative. Her attitude to both is bound up with her religion: she trusts God to care for her future, and she cannot understand why he allowed things to go wrong as they did. Her emotional landscape is a complex of many elements, present and past, including herself. Parts of the complex are positive, parts negative; parts goal conducive, parts goal obstructive; parts compatible with standards, parts not.

In that kind of situation, it is natural to associate dimensions with the individual events and actors, one of whom is the woman herself, as she is now, as she was, and as she hopes to be. It is natural to think of the situation in terms of a picture, or a series of pictures, where the events and actors are coloured to represent their emotional significance. (Colours are a visual way of expressing combinations of values on relevant dimensions that are associated with the events and actors.)

This may sound elaborate, but it is a simple practical requirement that coders need guidelines to specify which elements of a situation they should associate dimensional values with. The problem solves itself when one entity utterly dominates a person’s emotional life, like James’ bear. But everyday life is not as always as cohesive as that, and descriptive schemes need to have ways of dealing with the situations that everyday life throws up.

4.3 Dimensions in semantic memory

The last section considered how dimensions can be integrated into descriptions of individual emotional episodes. It is equally natural to integrate dimensions into what psychology has typically described as semantic (in contrast to episodic) types of representation. Two large bodies of work use that approach, one concerned with word meaning and one with attitude.

The work on word meaning stems from Osgood, who has already been mentioned. Research in that tradition showed that a large part of the lexicon had substantial affective connotations, which could be captured reasonably well in terms of E-P-A components ('affect' is used here deliberately, in the sense of section 1). So, for instance, motherhood and apple pie are high on evaluation and low on potency; sin and death are the opposite.

That approach to semantics was sidelined with the emergence of classical AI, which conceptualised semantics in ways that were much easier to use, and in particular, to use in inference and the control of action. However, it should be clear by this stage that affective components are far from irrelevant to inference and the control of action. Because motherhood and apple pie are high on evaluation and low on potency, there is a strong implication that one ought to protect them, which politicians notoriously but effectively exploit.

The general point here is that dimensional representations of semantic affect are not particularly useful in isolation, but if they can be coupled to appropriate schemes for inference and control of action, then they become very seriously relevant to some areas where humans and machines are likely to interact. Personal trainers need to know that 'slog' is negatively toned, in contrast to 'achievement'.

The research on attitude has a great deal in common with the Osgood tradition. It deals with enduring evaluations on a small number of dimensions, which it is assumed can be elicited using rating scales; but the evaluations are individual rather than embedded in the common language. The dimensions involve evaluation in a broad sense, but there may be more than one evaluative dimension.

After much debate, it seems clear that attitude as defined in this tradition can influence behaviour, though only in specific circumstances; and that its influence is linked to the fact that it involves affect, not simply cognitive evaluation. Again, dealing with attitude in this sense is relevant to various areas where humans and machines are likely to interact. Machines that do not understand their user's attitudes towards paperclips and perky dogs are paradigms of emotional incompetence.

5 Processes in emotional life

Emotion is strikingly dynamic, and it is quite paradoxical that descriptions of it are so often static – couched in terms of what seem to be steady states whose only time-related property is being off or on. Actually, trying to describe a static state makes it seem less paradoxical – it simply is extraordinarily difficult to describe even that satisfyingly, let alone a process that runs through a sequence of them. Nevertheless, it makes sense to try. This section tries by drawing up a list of what seem to be the most obvious kinds of process in emotional life.

It is part of psychology's ancient history that thinking in terms of process brings changes of perspective with it. That was a key part of the famous shift from structuralism (static) to functionalism (dynamic) in the late nineteenth century. So the list is not simply about covering topics that are missing from earlier sections – it is also about seeing topics that were covered earlier from a different perspective.

Part of the shift in perspective comes about because emotional processes are fundamentally interactive. They are about reacting to an environmental event, reaching out to other people, and so on; and their trajectory depends fundamentally on the interactions. As a very basic way of reflecting that, the descriptions given here have two parts, which will be called egocentric and allocentric. The egocentric part refers to processes within the subject, the allocentric to their external inputs and expressions.

Different parts of the description are developed to very different levels. Obviously the parts that have been thought through to some extent are models for the way the sketchier ones might be developed. On the other hand, none of the material is very well developed, simply because the literature is mainly oriented to states rather than processes. The main point is to provide a framework that brings issues into view and provides a context that may help people to work on them in the longer term, rather than being pulled relentlessly back into working with state-based models.

It will also be obvious to people with a background in psychology that the approach has links to JJ Gibson's 'ecological psychology' (Gibson, 1979). The links are there at various levels. It emphasises everyday experience rather than supposedly pure cases identified by theory and created in the lab to match the theory; it tries to find a middle ground between standard vocabularies that are either purely mentalistic or purely physicalistic, and that are therefore awkward to connect; and it treats action as a central issue, not something to be dealt with later when the static case is fully understood.

On the other hand, the approach makes no claim to be philosophically radical or deep, or to adhere to principles laid down by Gibson. It also tries to avoid the classic pitfall of Gibsonian theorising – focusing too much on the elegantly chosen phrase, and too little on the kind of mechanism that would do what the phrase seems to describe. The aim is always to describe in ways that have a real bearing on the construction of a system capable of engaging usefully with human emotional life. Of course, aim and success are entirely different matters.

There is no immediately obvious rationale for choosing which processes in emotional life to discuss first, because they are not nested in any straightforward way. However, some starting points allow issues to be introduced more gradually than others. On that basis, it is convenient to start with processes that hinge on representations of an emotionally charged object. 'Object' is used on the sense of section 2. It may refer literally to an object, or to a person, or to a situation, or to a sequence of events.

Figure 8 sums up the processes that will be considered. The point is to convey that the descriptions in the rest of the section are not chosen at random. They are an attempt to pull together the main kinds of process that emotion-oriented computing has grounds to be interested in. Correspondingly, it makes perfect sense for people looking at it to suggest where it should be expanded or restructured to give better coverage of emotional life.

Figure 8 : an overview of key processes in emotional life

Egocentric	Allocentric
Having an emotionally charged object representation <ul style="list-style-type: none"> ▪ specific representations that carry an emotional orientation ▪ general representations that carry an emotional imperative ▪ being active and being dormant 	Conveying that a particular topic is charged Registering that an issue is emotionally charged for someone else Expressing and interpreting emotion according to context
Activating an emotionally charged object representation <ul style="list-style-type: none"> ▪ in response to the immediate surroundings ▪ sequence of appraisal checks ▪ temporal profiles: sudden & irruptive, gradually arising, intermittent emotion ▪ landscape and mindscape 	Evolution of facial expression Facilitating or inhibiting emotional activation
Cognitive representations of emotional life <ul style="list-style-type: none"> ▪ assigning emotion words to own state ▪ perceiving emotional conditions in self ▪ Remembering emotional states in self ▪ Formulating emotional narratives for self ▪ Imagining emotional states <p style="text-align: right;">— affective</p>	...and states of others ...and others (including having empathy) ...and others ...and others ...enacting emotional states loops —
Forming an emotionally charged representation <ul style="list-style-type: none"> ▪ inheriting an emotionally charged representation ▪ classical conditioning ▪ transfer (Tomkins) ▪ attachment (Bowlby) ▪ empathy 	Conveying that the representation is being formed Swaying people emotionally
Perceiving, thinking emotionally Exerting rational control of emotional tendencies in perception, thought	Talking or acting emotionally Exerting rational control of emotional tendencies in talk, and action

- Suppressing emotion	- Suppressing signs of emotion
Having a personal emotional setting or predisposition	Conveying the global setting

5.1 Having an emotionally charged object representation

Perhaps not all emotion-related processes have objects (mood is conceptually defined as objectless, though that is no guarantee that such states exist). However, most do have objects. Talking about these objects in the context of process, it is useful to build on the ideas introduced in section 4.2, and to think of the objects of emotion as representations which are coloured or charged. Hence the first processes to be considered are the processes associated with having an emotionally charged representation of an object (or a person, or a situation, or a sequence of events).

It is standard to ask questions about objectless emotions, but starting with objects draws attention to the mirror question: can emotionally charged representations of objects be subjectless? Intuitively, it seems likely that emotionality implies a connection with the subject – not simply codings that identify something as good or bad, powerful or powerless, but also links that give those qualities a personal significance. The issue is not purely a matter of words: it has a bearing on the way one might imagine modelling emotion-related representations. It is not easy to see ways of addressing it empirically, though.

5.1.1 Egocentric processes

The egocentric processes associated with having an emotionally charged representation seem to divide into two levels, long term and short-term.

For the long term, it seems to be worth separating two extreme types of emotionally charged representation that people may have.

One is a specific representation that carries an emotional orientation, such as (in my case) an island that I particularly love, or a politician that I particularly detest. When something calls either one to mind, it is likely to bring strong feelings with it, but not usually any particular action – not least because the object of the feelings is not usually there. Specific representations that trigger action rather than feeling do exist (the rare person one would spontaneously hit or hug at first sight), but that seems to be the exception.

The second is a general representation that carries an emotional imperative. Everyone must have experiences of confronting situations that are quite novel in their details, but that when they happened, compelled intense emotional reactions. Examples in my own experience include (negatively) being stuck halfway up a cliff that was much trickier than I thought; and (positively) seeing island after island (as it seemed) unimaginably away from a high mountaintop. From the standpoint of cognitive psychology, it is natural to infer that both experiences must have activated some kind of long term representation, but one that deals with a class of objects or situations rather than an individual instances (Cowie 2002 reviews ideas about the emotional power of landscapes).

For the short term, the obvious distinction is between being active and being dormant. Being dormant means that a representation is stored in some way, but not materially influencing thought or behaviour. At any one time, most of the emotionally charged representations that a person has are in that state. Being active means that the representation is exerting an influence on the person's thoughts, feelings, or actions.

There seem to be in-between states. For example, preoccupations seem to be continually probing for opportunities to become fully active.

5.1.2 Allocentric processes

The obvious allocentric counterparts of having a charged representation are conveying that a particular topic is charged, and registering that an issue is emotionally charged for someone else.

It is an important matter for human-human communication to understand the topics that are emotionally charged for the different parties – topics that should be avoided because they are potentially explosive, or that have the potential to create bonds. It is reasonable to assume that similar kinds of issue will tend to arise in human-machine interaction as it becomes less restricted. Humans are particularly aware of the cases that differentiate people, but of course, they are marginal relative to the very large body of topics on which people at least within a single culture are likely to feel similarly. Machines with the ability to interact on a wide range of topics, but without any sense of those shared feelings, would be expected provoke problem reactions on a regular basis.

Despite the seeming importance of the issue, the literature on it seems to be very limited. There are techniques in marketing research that seem to be concerned with roughly this problem – finding out who feels what about particular brands of car or types of music. But it is not easy to find material on the way people recognise sensitive spots in the course of normal interaction and observation (after all, actions speak louder than words), or build up a picture of another person's sensitivities. Of course there are simple indicators like direct association between a topic and overt signs of emotion, but intuitively others seem to be just as important – avoiding topics, or raising them for no obvious reason; showing emotion after a charged topic has been safely left behind; aligning with people who take a particular stance on a topic, or drawing away from them; using particular kinds of language; resisting arguments that appear to be sound when they affect particular issues; and so on.

These issues do not come to mind easily if one starts by thinking about states, but that does not mean they are esoteric. They are bread and butter issues that have to be addressed if a machine is to be effective as a teacher (which bits of mathematics strike terror into the pupil, which does he or she find fun or fascinating?), or a health advisor (if the user hates the feel of water, stop recommending swimming); or a holiday advisor (are crowded, baking beaches his or her dream of heaven or hell?).

An issue that affects the allocentric side of most processes is how the expression and interpretation of emotion depend on context, social and relative to ongoing tasks. For instance, Hansen (2006) has studied expressions of fear in soldiers appearing before a military tribunal. It seems unlikely that the signs given there will be the same those given by a child who sees a frightening face at the window during family tea. Conversely, learning about the relevant representations will be very different in the two cases. The issue applies throughout the section on process, but it will only be raised here.

5.2 Activating an emotionally charged object representation

The last section separated having an active representation from having a dormant one, but did not discuss the transition process. This section takes up that issue.

5.2.1 Egocentric processes

One of the most studied processes in emotional life is activation of an emotional representation in response to the immediate surroundings. James took the lead with his proposal that action and visceral reaction preceded changes in consciousness. Arnold filled in a gap by proposing that appraisal processes governed the action and reaction, and Zajonc reasserted the late arrival of cognitive elements.

The best developed contemporary account is due to Scherer's group. It describes a sequence of 'stimulus evaluation checks' which makes sense logically, and which fits data collected by the group. Details are in Sander et al (2005), but briefly, they propose that checks occur in the following sequence:

- Relevance (including sub-checks for novelty, intrinsic pleasantness, and relevance to the subject's goals & 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)
- Normative significance relates to the way outcomes relate to one's own values and to society's

Activation very probably does follow a pattern of that general kind, but it seems most unlikely that it is the only pattern to occur. The point can be made from two directions. First, the pattern fits one of the three temporal envelopes noted by Wundt (that is, sudden & irruptive), but not the others

(gradually arising and intermittent emotions). Second, it applies to cases where the object is, so to speak, part of the landscape; but not to cases where it is part of the mindscape. The term 'mindscape' is used to cover both cases where the relevant experiences are private (anticipating or remembering in isolation) and those where they are socially evoked (most obviously as part of a conversation).

It would seem intuitively that the two are linked. Sudden, irruptive emotions are generally prompted by 'landscape' events, though there may be exceptions (suddenly realising that one has completely missed an important appointment). Gradually arising emotion may be prompted by events in the landscape (irritation at poor service) or in the mindscape (sharing memories of a happy time). Intermittent emotions are often (in the term used earlier) established, which means that they are rooted in mindscape events, remembered or anticipated.

Another range of linked issues flows from the fact that 'activation' is (deliberately) a blanket term. There are deep differences between bringing to mind a politician's grin and swagger and grasping that the category 'one false move away from a lethal fall' applies here and now. These are important for a deep analysis of processes in emotional life, but beyond the scope of this paper.

The kind of analysis that Scherer's group has provided for sudden and irruptive activation seems not to be available for the other patterns. It may well be that the task is simply more difficult,

because more complex kinds of interaction are likely to occur in the more protracted processes. That is linked to the allocentric perspective.

5.2.2 Allocentric processes

One level of allocentric process is conveying that emotional activation is occurring. Again, Scherer's group has developed elegant analyses of the process for sudden irruptive emotions, showing that facial expression evolves through a sequence of qualitatively different states, each related to a particular stage of appraisal, rather than morphing uniformly from a neutral to a peak state and back.

Analysing the signs associated with the other types depends on appropriate databases. It seems likely that they will have to be naturalistic, since one of the hallmarks of acted data appears to be that transitions between states tend to be abrupt and holistic.

A major issue on the allocentric side is facilitating or inhibiting emotional activation in another person. As mentioned above, that applies particularly with gradually arising and intermittent profiles. Understanding how to change these trajectories is one of the most obvious goals for emotion-oriented computing. A useful marker of the state of progress is the crudeness of the induction techniques used in emotion research, and the limited range of emotional states that they can evoke reliably.

5.3 Cognitive representations of emotional life

This section is concerned with the processes involved in talking and thinking about emotional life. It starts with the tokens that are at the centre of the process, everyday emotion words.

The section deals with both egocentric and allocentric issues, but in a slightly different way from the previous ones. The reason is that cognition depends on a special kind of relationship between the two. It is particularly clear in the case of words, and so the discussion begins there.

5.3.1 Assigning emotion words

It is highly natural to regard emotion words as labels corresponding – rather crudely and approximately – to inner states. A substantial tradition in philosophy points out that such ideas are fundamentally misguided. The rules of use for words that refer to mental life must allow people other than the subject to apply them (otherwise they could not form part of a common language). Because of that, words must be rather complicated instruments whose rules of use are logically bound to both ego- and allocentric elements.

This section is based on a recent reply to an article in that tradition by Sabini and Silver (Cowie 2005). The article sets out analyses of selected emotion terms in by bringing together observations that Sabini and Silver assume a competent user of the language would be unlikely to dispute. The reply considers what kind of system their analysis suggests might be needed to assign emotion-related words as humans do. It identifies eight types of criterion that it seems necessary to consider in assigning emotion terms.

Source criteria The term source seems as good as any to describe objective events or situations which prompt an emotion. Sabini and Silver make it clear that the actual source is relevant to the choice of word. They consider which of two labels, anger and envy, should be applied to a person

who claims his friend bragged about being offered a place at Harvard. It matters whether a friend actually did brag about the offer, not just whether the first party thinks he did.

Appraisal criteria Sabini and Silver use concepts like a perception of insult which seem for all practical purposes to be appraisals. The appraisal needs to be considered as well as the source. The point is made using the example of Lord Jim from Conrad's novel who jumped ship when he believed that it was sinking (though it was actually not). His fear was driven by a mistaken appraisal, not the actual state of the ship.

Impulse criteria Impulse is relevant to assigning emotion words from both the subject's standpoint and an outsider's. From the subject's standpoint, impulse intrudes into consciousness, and that kind of intrusion is characteristic of emotion. From the outsider's standpoint, impulsive actions are a marker of emotion.

Involuntary signs This term is meant to cover elements like visible arousal and facial expression, which again are available to both first and third parties.

Action criteria Sabini and Silver cite specific types of voluntary action that are relevant to assigning an emotion word – taking revenge, fleeing, complaining.

Action evaluation Several of Sabini and Silver's examples highlight the fact that assigning an emotion term entails an evaluation of the first person's actions and reactions – particularly in cases such as jealousy, pride, and regret.

Character evaluation According to standard analyses of the words 'shame' and 'embarrassment', the difference is that 'shame' describes reactions to a real personal deficiency being exposed; whereas 'embarrassment' describes a situation where the person is somehow (perhaps unfairly) made to seem deficient. Hence the choice depends on evaluation of the real and apparent qualities of the first person.

Valenced feeling states These are highly relevant to assigning an emotion word from the first person standpoint, but not available (unlike most of the other criteria) to the outsider.

That kind of list is a first step because, as Sabini and Silver make clear, various other issues surround it. A critical one is that the different types of criterion play different roles in different cases. Some words can only be assigned when there exists a highly specific combination of impulses and feeling states (their discussion of shame includes a vivid description). For others, there is a slot which may be filled by various types of state – in the case of regret, various types of negative feeling. One might suspect that envy and anger illustrated a counterpart – the attribution of envy implies rather specific types of action evaluation, the attribution of anger less so.

Cutting across these themes is the issue of perspectives. It has two aspects. The first is just that both first- and third-party perspectives need to be considered. The criteria that Sabini and Silver emphasise are interesting in that most of them are as relevant to a third party as to a first, though not necessarily in the same way. The second is that the same emotional experience may be described in different ways. For instance, Lord Jim's emotional response to a memory can legitimately be understood in terms of either shame, or remorse: it forms part of both patterns, and the ascriptions are not in competition. Sabini and Silver lay a great deal of emphasis on that point, arguing that the number of feeling states is considerably smaller than the number of emotion words: many emotion words are essentially contextualising the same feeling state in different ways.

A natural way to develop these ideas is to formulate scripts for emotion words based on the headings in the list. These would specify the events that are to be expected if the script applies, including information about what should be apparent from which perspective. An attraction of that format is that it points naturally towards the inferential character of emotion terms – if matters on which there is evidence conform well enough to a particular script, then one infers that other matters, on which evidence is missing, or which are in the future, will probably also conform to it.

It is because emotion terms support inferences that they are potentially useful to machines – they go beyond the information given, and make it possible to anticipate or understand. However, the list also highlights types of inference entailed in everyday terms that we might not be comfortable for a machine to make. Emotion terms entail evaluations, up to and including imputations of a permanent character flaw. It is galling enough when a human being applies the term ‘sulk’ to what we experience as justified resentment. A machine that took the same tack might expect to be smashed. The issue here is that entitlement is also part of the process of assigning an emotion word.

The general point is that the process of assigning everyday words is intricate, and should be taken seriously for what it is rather than being assimilated to an image which is appealing conceptually, but a misleading model of the way real words are put to use.

5.3.2 Perceiving emotional conditions

Emotional life involves at least two ways of perceiving emotional condition – perceiving one’s own emotional condition and perceiving someone else’s. Starting with the second makes it easier to identify the issues.

It is tempting to identify perceiving someone else’s emotional condition with attaching a state label. Attempts to label databases make it very clear that people perceive something much richer and less easily defined than that. In fact, figure 7 is to a large extent based on trying to capture the main things that people automatically try to grasp about another person’s emotional condition when they watch (and listen to) a recording.

Philosophers in particular have argued that one of the key resources used in perceiving someone else’s emotional condition is empathy – ‘standing in the other person’s shoes’. Optimists suppose that empathy like that gives our understanding of another person’s mind a kind of authority that is quite unlike ‘naïve theories’ in other fields (physics, biology, etc).

The idea is intriguing, but experience suggests that empathy is at least as error-prone as ‘naïve physics’. People systematically overestimate other people’s similarity to themselves (Wilhelm & Perrez 2004), and Goldie (2006) argues that that is almost inevitable. Perhaps empathy could be accurate if we could recreate all the conditions that bear on a person’s feelings at a given time. But in truth, we rarely have much of the relevant information; and if we had, we could not strip our own emotion-forming background out of our emotional machinery and replace it with the other person’s. Nor, for that matter, is the machinery that we use to anticipate our own emotional reactions anything like foolproof. We hate parties that we expect to love, and we enjoy recitals that we went to gritting our teeth.

Those issues are linked to perceiving one’s own emotional condition. The classical debate has been whether being in a particular emotional condition entails knowing it. James, Freud, Zajonc, and many others have argued for the answer ‘no’, and that is surely in line with the all-too-familiar experience of having to be told by a third party that one is not behaving quite rationally. That raises the question of what being emotionally self-aware is.

Certainly one model is that emotional self-awareness is simply opening shutters that may (but need not) close off certain areas of mental life from conscious awareness: and when the barriers are gone, the processes simply become part of conscious life. However, it seems more likely to be creating a representation of oneself at a different level. It may be that that level is symbolic, or it may be that it exploits the machinery of being emotional in an intermediate way, like a kind of empathy with oneself – perhaps rather as picture perception seems to exploit the mechanisms of seeing what is literally there.

These issues are certainly difficult in many ways. However, they are very real, and very relevant to emotion-oriented computing. For instance, one of the obvious applications of emotion-oriented computing is giving users feedback to help them stay within a range of emotionality that they regard as appropriate – that is, supplementing the process of representing emotion cognitively.

5.3.3 Recall processes

Intuitively it would seem that the process of remembering emotional episodes must be extremely complex and error-prone. Commonplace experiences illustrate the point – for instance, coming across something with old associations (a photograph, a place, a tune, a smell, a piece of writing) and finding that it brings back emotions from a region of one's past that one had not remembered in that way at all.

A lot of data confirms a wary view of memory for emotion. Data from Myrtek (2004) suggest that people feel a specific emotion over 10 times a day. But asked to recall an emotion the previous day, about 50% seem unable to (Scherer et al 2004). Forgetting is not only massive but selective: memory is loaded in favour of episodes where the emotion is particularly intense (Reisberg et al 1988) and overestimates how negative emotions were (Wilhelm, Schoebi & Perrez 2004). It is also a basic finding in memory research that if something has been categorised, one's memory of it tends towards a typical instance of the category. There is much more, but perhaps the point is made.

Part of the trickiness of memory is that it is not only concerned with recalling individual events – it is also concerned with formulating emotional narratives (Goldie 2001: Wilhelm, Schoebi & Perrez 2004). People are concerned to find patterns in their emotional experiences over a period, quite reasonably. But they are capable of distorting the evidence to fit the pattern, as much in memory as elsewhere.

There are two reasons for being interested in this material. First, researchers need to be wary of their unaided memories as sources of information about their emotional lives. Combining Myrtek's and Scherer's figures, it appears that 95% of emotional episodes has faded from memory by the next day; and what remains.

Second, one of the areas where emotion-oriented computing has potential is in giving people access to better records of their own emotions. 'Life albums' are not a new idea, but they could be a great deal more than toys.

5.3.4 Synthetic processes

It is a prominent feature of human life that people synthesise or simulate emotions in various ways. Perhaps the basis of the activity is empathy, which was discussed above, but it is enormously important practically – after all, the entertainment industry depends on it.

The pure egocentric form is recalling or imagining an emotional state. In the pure form, it is perhaps not a particularly powerful or useful phenomenon (though it is favoured by some teams working on affective computing precisely because it gives a credible approximation to affect per se).

A pure allocentric form is acting an emotion. It is easy to shrug off acting as a mere commercial phenomenon, of no scientific interest. But it very interesting that people have a certain basic capacity to project the signs of an emotion without feeling it in any depth, and equally interesting that certain mental disciplines (which actors learn) can result in signs that are more like signs of real emotion. It is quite fascinating that (as Shakespeare pointed out) people in some ways find the acted version more satisfying than samples drawn from real emotional episodes – from audiences in the Shakespearian theatre to emotion researchers. Perhaps the actors correspond to the ‘superstimuli’ that draw sticklebacks into mating efforts far beyond those they would make for an actual stickleback.

There is a whole complex of issues here for emotion-oriented computing to address. It is important to be clear how different acted data are from what happened in life. It is also important to know whether the best data to use are acted or natural. It is not obvious that superhero in an adventure game should be subject to all the mundane features of real human emotion. Sometimes, emotion-oriented computing may be better giving us the emotions we imagine than the ones we actually have.

Cutting across ego- and allocentric is a large body of processes where (in the vernacular) people work themselves into a state. Working into a state is usually a social phenomenon – at least, it is easiest when there is another party. Tam o’ Shanter’s wife is an exception, sitting at home

Gathering her brows like gathering storm
Nursing her wrath to keep it warm.

When there are multiple parties, Hook’s concept of an affective loop seems to come into play. Emotionally charged actions produce feedback that promotes more actions, and so on. Little seems to be known about the process – which kinds of action/feedback combinations amplify, and why.

These loop phenomena lend themselves to what has been called an ‘emotional gym’ (Cowie et al, 2003), where the person can ‘work out’ (if he or she so chooses). People enjoy emotional workouts without all the consequences that are liable to follow from a workout with a real human, and that too is both interesting scientifically and potentially important in terms of applications.

5.4 Forming an emotionally charged representation

It is obvious from sections 5.1 and 5.2 that there are questions about the ways emotionally charged representations are formed. The question has been postponed until this point because it does not seem very relevant to emotion-oriented computing, and it will be skimmed.

The basis of emotionally charged representation is presumably heredity. Situations of certain kinds evoke what look like emotional reactions from very early in life, and more ‘switch in’ in ways that suggest prespecified mechanisms gradually coming on stream.

One of the first mechanisms of learning to be analysed in depth, Pavlov’s classical conditioning, has strong links to emotion. Roughly speaking, it transfers the pre-existing emotional significance of one kind of object to another.

Tomkins (1992) proposed another mechanism, which he called ‘transfer’. He argued that affect is highly contagious because facial and voice patterns can elicit the complete pattern of affect – hence the core of one person’s affective stance can be transferred to another person.

Another major tradition in psychology proposes that individual emotional priorities and balances arise out of inability to meet innate needs – for instance, Bowlby argued that the need for secure attachments played a pivotal role in emotional development.

Naively, it is not obvious that these mechanisms can account for the variety and force of emotionally charged objects that shape a typical adult’s life. That may mean that different mechanisms need to be considered. It is not outrageous to suggest that mechanisms more like learning facts or establishing conclusions might be involved. However, that is not for this paper.

The obvious allocentric issues in this area are conveying that an emotionally charged representation is being formed, and deliberately swaying the way people feel emotionally about a person or topic. Emotion-oriented computing probably ought to consider these from the moral standpoint, at least. Shifting a person’s enduring valuation of the world is not a thing to be done lightly.

5.5 Perceiving, thinking and acting emotionally

There is a whole network of issues to be considered under this heading. Some have been studied in some depth, others hardly at all.

The literature on perception and emotion often does not distinguish between perceptual processes that evoke emotional responses (at some level) and perceptual processes that occur when someone is already in an emotional state.

The appraisal literature deals with perceptual processes that evoke large-scale systemic responses. Another literature deals with the more limited responses associated with seeing an emotionally charged picture (sometimes evoking an emotion, e.g. a spider, sometimes expressing it, e.g. a frightened face). Taylor and Fragopanagos (2005) provide a summary, with particular reference to the central role that attention plays in the phenomenon.

A well-known practical example 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). Laney et al (2004) argue that the effect is specific to perceptual processes that evoke emotional responses, and not to perceptual processes that occur when someone is already in an emotional state.

There is a substantial literature on the way one emotional state, anxiety, affects perceptual and related process that occur within it – attentional control, depth of processing, and speed of processing. Eysenck and his team have developed a theoretical analysis of the issues surrounding these effects (1997, 2007). Other emotions seem not to be covered in the same way, and it is an interesting issue whether Eysenck’s work is a model that can be extended. It should not be taken for granted, because the core of Eysenck’s analysis is that emotion-driven through processes consume critical space in working memory: that may well be specific to anxiety.

Positive affect has generally been studied in relation to more central processes. It promotes recall of material remembered under positive conditions, which tends to be useful because it is extensive and well-organised. It fosters flexible and creative thought, and can speed decision-making. It affects risk-taking, not necessarily in obvious ways (Isen 1999).

A famous case in the psychology of emotion illustrates the intuition that emotion is central not only to perception and cognition, but to action. Phineas Gage was the victim of a bizarre accident in 1848. Damasio famously argued that the key damage must have been to ventromedial prefrontal cortex, which allows emotional values to be integrated into planning. He recovered, but with irreparable damage which contemporaries noted from the way he acted:

He was now fitful, irreverent, and grossly profane, showing little deference for his fellows. He was also impatient and obstinate, yet capricious and vacillating, unable to settle on any of the plans he devised for future action.

<http://www.deakin.edu.au/hbs/psychology/gagepage/> :

The details of the case are highly debatable (McMillan 2000), but it does at least demonstrate how plausible people find the idea that emotion can have powerful effects on action. The law has accepted as much for centuries by giving specific recognition to ‘crimes of passion’ (*crimes passionnel* in French, *Affekthandlung* in German).

The idea that emotion can precipitate action has an equally familiar mirror image, also prominent in the interpretations of the Phineas Gage story, which is that a major part of emotion-cognition interaction consists of cognition exerting control over impulses precipitated by emotion (Wagar & Thagard 2004). These are a curious crossover between ego- and allocentric – exerting rational control of emotional tendencies in talk, and action in one sense affects the allocentric domain, and in another the whole point is that nothing untoward enters it.

This is a very brief glimpse over a large, complex area. The main point for emotion-oriented computing is to convey a little of the scale and practical importance of the phenomena that are in the literature, let alone those that are not formally documented (which many are not, because they are hard to control). Understanding these processes is one of the keys to identifying applications for emotion-oriented computing.

5.6 Having a personal emotional setting or predisposition

This heading is effectively a process-oriented way of describing issues that take centre stage in standard state-oriented approaches. Raising them at the end may help to suggest how much of emotional life does not hinge on the issues that those approaches bring to the fore.

In this context, it is natural to couch the question of what personal settings (long- or short-term) are in terms of the ways that they might link in to descriptions of the various processes that have been sketched already. It is tempting to tag them ‘ghosts in the machine’, and deny that they have any material influence in anything. That would be sophistry.

People know from their own experience that they sometimes have global biases, and will (for instance) see problems at any opportunity. They also know that something like the synchrony described in Scherer’s definition of emergent emotion is important in the domain of emotion. At certain points, elements ‘click’ and enter qualitatively different configurations. It is not clear to what extent these qualitative transitions are a function of the control mechanisms, to what extent of mechanisms for interpreting signs (one’s own or someone else’s), and to what extent from interactions. Letting a clear sign of anger escape changes the situation for the giver and the receiver.

From a process perspective, these are real and interesting issues. What is not obvious is how much attention has to be paid to them to achieve models that work to an acceptable approximation.

6 Conclusion

The theme of this paper is that emotional life is an extremely large, extremely complex tangle of phenomena, capable of being looked at usefully in several different ways. Some disciplines will certainly find that outlook unpalatable, and will blame psychologists for failing to deliver a solution on one page.

Unfortunately, the problem lies not in psychologists specifically, but in human beings generally. Cognition is very much the easiest part of human life to formulate clear theories about, and it is still very partially understood. Emotion is intrinsically much harder to square with a logical explanatory framework, and much less time has been devoted to trying. The wise response is to accept that for the foreseeable future, all maps will be approximate, and it is as well to be familiar with a variety.

7 References

Alleydog <http://www.alleydog.com/glossary/psychology-glossary.cfm>

Arnold MB (1960) *Emotion and personality I* Columbia University Press

Averill, J.R., 1980. A constructivist view of emotion. In: Plutchik, R., Kellerman, H. (Eds.), *Emotion: Theory, Research and Experience*, Vol. 1. Academic Press, New York, pp. 305–339

Ben Ze'ev, A. (2000) 'I Only Have Eyes For You': The Partiality of Positive Emotions *Journal for the Theory of Social Behaviour* 30 (3), 341–351

Cowie, R (2002) Charged experiences of natural environments *Current Psychology* **21**, 133-143

Cowie R and Cornelius R. (2003) Describing the Emotional States that are Expressed in Speech *Speech Communication* **40**, 5-32

Devillers, L., Vidrascu, L., Lamel, L. (2005) Challenges in real-life emotion annotation and machine learning based detection *Neural Networks* **18**(4) 407 - 422

Douglas-Cowie, E., Devillers, L., Martin, J-C., Cowie, R., Savvidou, S., Abrilian, S., Cox C. (2005). "Multimodal Databases of Everyday Emotion: Facing up to Complexity", *Interspeech 2005*.

Eysenck M. (1997) *Anxiety and Cognition: A Unified Theory* The Psychology Press

Eysenck, M.W., Derakshan, N., Santos, R., Calvo, M.G. (2007) Anxiety and cognitive performance: Attentional control theory *Emotion* 7(2) 336-353

Feldman Barrett, L. (2006) Solving the Emotion Paradox: Categorization and the Experience of Emotion *Personality and Social Psychology Review*, Vol. 10, No. 1, 20–46.

Frijda, N (1986) *The Emotions* Cambridge: Cambridge University Press

Gibson, JJ (1979) *The Ecological Approach to Visual Perception* Boston: Houghton Mifflin

Goldie, P (2001) *The Emotions: A Philosophical Exploration*. Oxford University Press Oxford.

Hilgard (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 *Medicine Meets Virtual Reality: The Convergence of Physical and Informational Technologies Options for a New Era in Healthcare*. JD Westwood, HM Hoffman, RA Robb, D Stredney. (eds). pp. 148-153. Amsterdam: IOS Press

Isen, A.M. (1998). On the relationship between affect and creative problem solving. In S. Russ (Ed.), *Affect, creative experience, and psychological adjustment* (pp. 3-17). Ann Arbor, MI: Braun-Brumfield.

James, W (1884) What is an Emotion? *Mind*. 9, 188-205

Kahneman, D. (1995). Varieties of counterfactual thinking. In N.J. Roese & J. M. Olson (Eds.), *What might have been: The social psychology of counterfactual thinking* (pp. 375-396). Mahwah, NJ: Erlbaum.

Laney C, Campbell HV, Heuer F, Reisberg D. (2004) Memory for thematically arousing events *Memory and Cognition* 32, 1149-59.

McMillan, M. (2000) *An Odd Kind of Fame: Stories of Phineas Gage*. Cambridge, MA: MIT Press

Myrtek (2004) *Heart and emotion: Ambulatory Monitoring Studies in Everyday Life* Gottingen: Hogrefe & Huber

Ortony A (2002) On making believable emotional agents believable. In: *Emotions in humans and artifacts* (Trapp R, Petta P, Payr S, eds), p 189. Cambridge, MA: MIT Press

Osgood, C., Suci, G., Tannenbaum (1967) *The Measurement of Meaning* Urbana: University of Illinois Press

Panksepp, J (2003) At the interface of the affective, behavioral, and cognitive neurosciences: Decoding the emotional feelings of the brain *Brain and Cognition* 52 4-14

- Rapaport, D. (1953). On the Psycho-Analytic Theory of Affects *International Journal of Psycho-Analysis*, 34:177-198
- Reisberg, D, Heuer, F., McLean, J & O'Shaughnessy, M (1988) The quantity, not the quality, of affect predicts memory vividness *Bulletin of the Psychonomic Society* 26, 100-103.
- Rosch, Eleanor (1978): Principles of categorization. In: Rosch, Eleanor, Lloyd, Barbara B. (Eds.) *Cognition and Categorization*. Lawrence Erlbaum Associat
- Russell, J. A.. & Fehr, B. (1994). Fuzzy concepts in a fuzzy hierarchy: Varieties of anger. *Journal of Personality and Social Psychology*, **67**, 186-205
- Russell, J A (2003) Core Affect and the Psychological Construction of Emotion *Psychological Review* 110, No. 1, 145–172
- Russell, J. A.. & Feldman Barrett L (1999) Core Affect, Prototypical Emotional Episodes, and Other Things Called *Emotion*: Dissecting the Elephant *Journal of Personality and Social Psychology*, **76**, 805-819.
- Russell, JA (1991) In Defense of a Prototype Approach to Emotion Concepts *Journal of Personality and Social Psychology* 10, 37-47
- Sabini J & Silver M (2005) Why emotion names and experiences don't neatly pair *Psychological Inquiry*. 16, 11-48
- Scherer, K.R., 1999. Appraisal theory. In: Dalglish, T., Power, 2321 M. (Eds.), *Handbook of Cognition and Emotion*. John Wiley, New York, pp. 637–663.
- Scherer, K, Wrانik, T, Tran, V, & Scherer, U (2004) Emotions in everyday life *Social Science Information* 43, No. 4, 499-570
- Schlosberg, H.(1954). Three Dimensions of Emotion *Psychological Review* 61 (2) 81-88
- Sternberg, R.J., (1988) Triangulating love. In: RJ Sternberg, ML Barnes *The Psychology of Love* Yale UP, Yale
- Stocker, Michael and Hegeman, Elizabeth 1992. *Valuing Emotions*. Cambridge: Cambridge University Press
- Taylor, J.T. and Fragopanagos, N. (2005) The interaction of attention and emotion *Neural Networks* **18**(4) 353 – 369

Tomkins S.S. (1964) *Affect Imagery Consciousness: Volume I*. London: Tavistock.

Tomkins S.S. *Affect Imagery Consciousness: Volume 3. The Negative Affects: Fear And Anger*. New York: Springer, 1991

Wagar, B M. & Thagard, P. (2004) Spiking Phineas Gage: A Neurocomputational Theory of Cognitive–Affective Integration in Decision Making *Psychological Review* 111, 67–79

Wilhelm & Perrez (2004) How is my partner feeling in different daily-life settings? *Social Indicators Research* 67, 183-246.

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". *Social Science Information*, 43(4), 647-665

Wundt, W. (1897). *Outlines of psychology*. Translated by Charles Hubbard Judd <http://psychclassics.yorku.ca/Wundt/Outlines/>

Towards a consensus

There is a widespread sense that there are aspects of human life, related to emotion or affect, that Information Technology can and should engage with more systematically. However, it is not easy to give specific descriptions of the domain that Information Technology should be addressing.

At a meeting sponsored by HUMAINE at the University of Haifa in May 2007, a considered description of the domain was agreed by an international group of philosophers, psychologists, neuroscientists, and experts in emotion-oriented / affective computing. It is given below.

The declaration of Haifa

Information technology has been increasingly concerned with a domain that seems important intuitively, but which is difficult to describe explicitly. The group regarded the following as a reasonable way to summarise the domain.

It incorporates the aspects of life that are not pure cognition, appetite or will, including emotions, moods, stances, bonds, and altered states of control, arousal and engagement.

These aspects involve units on several levels, episodic, enduring, and pervasive; and involve some or all of the following components

- Global feelings and emotional colouring,
 - represented in terms of affective dimensions or appraisals;
- Distinctive action tendencies, modes of cognition & expression;
- Distinctive kinds of impression of relevant people, self & others;
- Rich connections to events and people past, present and foreseen

There are distinctive everyday terms to label some of the states and aspects involved, but not all.

A summary that reflects a large part of this is

‘felt appraisals, states, connections, expressions and tendencies’.

The domain is intimately linked to values and moral judgment. It is dependent on development and culture, and related phenomena are found in other species. It is more or less what Hume meant by ‘the heart’.