

Researching Emotions in Speech

Gordon McIntyre¹ and Roland Göcke

¹Research School of Information Sciences and Engineering, Australian National University
Vision Science, Technology and Application, NICTA Canberra Laboratory
¹gordon.mcintyre@anu.edu.au, roland.goecke@nicta.com.au

Abstract

Many studies into affective computing cover similar ground or overlap, pointing to much duplication of effort. Two problem areas exist in the research. The first relates to the way in which emotions are defined and described. The second is in the techniques used to elicit samples of speech. Further, few studies are based on natural speech, ignoring social, contextual, cultural and agent characteristics and few take account of the affective state of the receiver. This paper presents a survey of the problem areas and proposes 1) a comprehensive model of affective communication; and 2) a set of ontologies of affective computing.

1. Introduction

The December 2005 edition of *New Scientist* published an article titled, "Tell Laura I love her..." (Bennett 2005), featuring an insight into the work of esteemed, affective-computing pioneer, Ros Picard. The article, whilst positive, did air the view of one critic who complained,

"...that there is no valid theory, or even commonly accepted definition of 'emotion' to guide the work. These are smart people and they're producing useful things, but I think that's despite the framework they work in, not because of it...If the banner concept of your research is something that you can't even define, ...then you run the risk of producing things that do not advance human knowledge, but can only be described as 'cute' a digital motivator for a fitness programme, for example."

Overall, the article presented Picard and affective computing in a positive light. However, the criticism is understandable. Science has many definitions. According to Wikipedia, notions of *systematic*, *verifiable means* and *acquiring knowledge through empiricism*, are important. It is quite reasonable then for scientists from other areas of interest to expect that published results are built on a foundation that is systematic and verifiable.

Research in the field of affective computing is not consistent and, arguably, not verifiable. There is a lack of consistency in the way in which emotion is defined, measured and categorised. This in turn makes verification and reuse of results difficult.

This paper presents a survey of two of the major areas of contention in affective computing and proposes a more complete model that could be used to progress research into emotions in speech.

The next section explains the motivation behind this study. Section 3 provides some background to some fundamental problems in defining, categorizing and measuring emotions. Section 4 summarises the most common techniques for eliciting data for research. Section 5 presents a new, more comprehensive model and a set of ontologies

for use when researching the field. Section 6 concludes the paper.

The terms affect, affective state and emotion, although not strictly the same, are used interchangeably in this paper.

2. Motivation

The human computer interface is, for other than basic I/O operations, limited without some level of incorporation of the affective human state. An interface that responds and adapts to emotions can be more satisfying, much more productive and can provide many opportunities. One of the most natural interfaces is speech and we know that evidence of emotions are often present in speech. Emotion recognition in speech by computers could provide many benefits.

However, to accelerate progress, there needs to be more consistent modelling and reporting of affective computing.

3. Problem 1 - describing emotions

There is great difficulty in defining a term that is both ontological and overloaded with meaning.

Emotion theory is a large amalgam of approaches to the study of emotion, mostly based on cognitive psychology but with contributions from learning theory, physiological psychology, clinical psychology, and other disciplines including philosophy.

There are many reviews on emotional classification, e.g. (Cowie and Cornelius 2003). Two approaches to describing emotions seem to be dominant. The first is to define categories and is the more prevalent approach. The second approach is through the use of dimensions. Another approach, discussed briefly, is through the use of appraisal theory.

3.0.1. The category approach

The most intuitive and common approach in emotion recognition is the use of discrete prototypes. The "big six", which includes fear, anger, happiness, sadness, surprise, and disgust is quite popular (Cornelius 1996; Ekman 1999). These are regarded as "full-blown" emotions (Scherer 1999) and are evolutionary, developmental and cross-cultural in nature (Ekman and Oster 1982). However, there are many alternative groupings both across disciplines

and within disciplines. Some studies concentrate on only one or two select categories. Other employ schemes using more than twenty emotional archetypes.

One of the difficulties in comparing studies into emotions in research is that the choice of categories for a study varies and will often depend on an application that the researcher has in mind. If the focus of the research is to understand full-blown emotions then the big-six or a subset might be adequate. However, if the aim is to study the less dramatic emotional states in everyday life, with all the shades and nuances that we know distinguish them, then the choice of categories is much more difficult (Schröder and Cowie 2005).

Thus, there are many problems with the category approach such as:

- no agreed number of categories or definitions
- the large number of words with overlapping meanings
- lack of consistency of words across languages
- inconsistency in the use of categories in research

3.0.2. The dimensional approach

Another way to label affective state is to use multiple dimensions. Instead of choosing discrete labels, one or more continuous scales, such as pleasant/unpleasant, attention/rejection or simple/complicated are used. Two common scales are valence (negative/positive) and arousal (calm/excited). Valence describes the degree of positivity or negativity of an emotion or mood; arousal describes the level of activation or emotional excitement. Sometimes a third dimension, control or attention, is used to address the internal or external source of emotion. A software application has been developed at (Cowie 2000) to assist in the continuous tracking of emotional state on two dimensions. The package, called FEELTRACE, has been used in a number of studies.

3.0.3. The appraisal approach

Scherer has studied the assessment process in humans and suggests that people affectively appraise events with respect to novelty, intrinsic pleasantness, goal/need significance, coping, and norm/self compatibility (Scherer 1999). It is not yet clear how to implement this approach in practice.

3.1. Discussion

Both the categorical and dimensional approaches suffer from being highly subjective. Recent studies have shown that this is not only because of complexity but also because of the differences in the efficiency of listeners' physiology and the fact that the listener's own affective state influences their judgment of the speaker's affective state. Hence, there is a need for a model that includes listener attributes.

Analysis of Belfast Naturalistic Database has shown that full-blown emotions occur more rarely in natural speech than was thought (Cowie and Cornelius 2003). A study by the Human-Machine Interaction Network on Emotion (HUMAINE) group proposes that emotions be dichotomised into *episodic* and *pervasive* categories.

Episodic emotions are more like the traditional set of "full-blown" emotions or what Scherer would describe as emotions. *Pervasive* are the everyday emotions such as grief/sorrow, sarcasm/irony and surprise/astonishment. Regardless of the method employed to describe emotions, this greatly complicates the task of describing emotional content in speech samples.

A recent study by Devillers et al. has included labelling of *blended* and *secondary* emotions in a corpus of medical emergency call centre dialogues as well as including task-specific context annotation to one of the corpora (Devillers, Vidrascu, and Lamel 2005). Introducing the concept of multiple, concurrent emotions further compounds the difficulty of describing emotional content in speech samples.

4. Problem 2 - eliciting emotional speech

Most studies into affective communication begin with the collection of audio and/or video communication samples. The topic of collection of emotional speech has been well covered by other reviews (Scherer 2003; Cowie and Cornelius 2003) (Cowie et al. 2005) so it is only briefly covered here.

4.1. Naturally occurring speech

To date, call centre recordings, recordings of pilot conversations, and news readings have provided sensible sources of data to research emotions in speech. These types of samples have the highest ecological validity. However, aside from the copyright and privacy issues, it is very difficult to construct a database of emotional speech from this sort of naturally occurring emotional data sources. In audio samples there are the complications of background noise and overlapping utterances. In video there are difficulties in detecting moving faces and facial expressions. A further complication is the suppression of emotional behaviour as in the case where the speaker is aware of being recorded.

4.2. Induced Emotional Speech

Velten introduced the technique of getting subjects to read emotive texts and passages to induce emotional states (Velten 1968). Other techniques include the use of Wizard of Oz setups where, for example, a dialog between a human and a computer is controlled without the knowledge of the human. This method has the benefit of providing a degree of control over the dialogue and can simulate a natural setting.

The principal shortcoming of these methods is that the response to stimuli may induce different emotional states in different people.

4.3. Acted Emotional Speech

The method, most commonly used, is to use actors to portray emotions. This technique provides for a lot of experimental control over a range of emotions and like the previous method provides for a degree of control over the ambient conditions.

The problem is that acted speech elicits how emotions should be portrayed, not necessarily how they are portrayed. The other serious drawback is that acted emotions are unlikely to derive from emotions in the way that Scherer

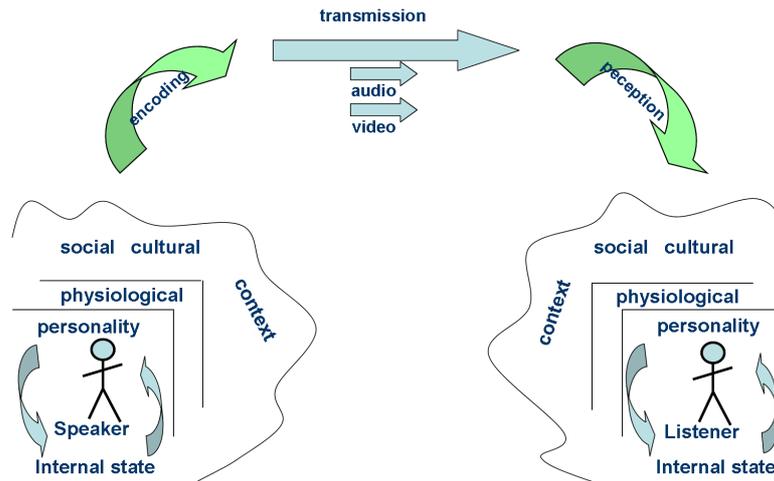


Figure 1: A model of affective communication

et al. describe them, i.e. “episodes of massive, synchronised recruitment of mental and somatic resources to adapt or cope with a stimulus event subjectively appraised as being highly pertinent to the needs, goals and values of the individual (Scherer 2004)”.

4.4. Discussion

Relatively little research into affect has been based on natural speech. In many cases, the approach to affect recognition has simply been an extension of automatic speech recognition (ASR), i.e. taking a corpus of acted speech, then annotating sequences containing affect within the corpus. In the case of automatic recognition of *episodic* emotions, this approach is plausible, based on the assumption that clear-cut bursts of episodic emotion will look and sound somewhat similar in most contexts (Cowie et al. 2005). However, recognition of *pervasive* emotions present a much greater challenge and, intuitively, one would think that awareness of personal and contextual information needs to be integrated into the recognition process. Richard Stibbard (Stibbard 2001) who undertook the, somewhat troubled, Leeds Emotion in Speech Project reported,

“The use of genuine spoken data has revealed that the type of data commonly used gives an oversimplified picture of emotional expression. It is recommended that future work cease looking for stable phonetic correlates of emotions and look instead at dynamic speech features, that the classification of the emotions be reconsidered, and that more account be taken of the complex relationship between eliciting event, emotion, and expression.”

5. The proposed solution

The proposed solution consists of 1) a generic model of affective communication; and 2) a set of *ontologies*. An ontology is a statement of concepts which facilitates the specification of an agreed vocabulary within a domain of

interest. Agents that wish to operate within the domain of interest commit to the vocabulary and its terms.

The model and related ontologies are intended to be used in conjunction to describe 1) affective communication concepts; 2) affective computing research; and 3) affective computing resources.

5.1. The model

Figure 1 presents the base model. The first point to note is that it includes speaker and listener, more in keeping with the Brunswikian lens model as proposed by Scherer (Scherer 2003). The reason for modelling attributes of both speaker and listener is that the listener’s cultural and social presentation vis-à-vis the speaker may also influence emotional content. The second point to note is that it includes a number of factors that influence the expression of affective in speech. Each of these factors and the motivation for their inclusion is briefly discussed. More attention is given to context as this is seen as a much neglected modulating factor in the study of automatic emotion recognition and synthesis.

5.1.1. Context

Context is linked to modality and emotion is strongly multi-modal (Cowie et al. 2005) in the way that certain emotions manifest themselves favouring one modality over the other. Physiological measurements change depending on whether a subject is sedentary or mobile. A stressful context such as an emergency hot-line, air-traffic control or a war zone is likely to yield more examples of affect than everyday conversation.

Stibbard (Stibbard 2001), has recommended,

“...the expansion of the data collected to include relevant non-phonetic factors including contextual and inter-personal information.”

His findings underline the fact that most studies take place in an artificial environment, ignoring social, cultural, contextual and personality aspects which, in natural situations, are major factors modulating speech and affect presenta-

Modulating factors			Production and detection factors		
Cultural	Social	Context	Agent Characteristics	Physiological	Internal State
Speaker's vis-à-vis listener's age and gender	Education	Group situations	Extrovert/Introvert	Voice quality	Recent events, eg lottery wins, losses
Language	Familiarity/rapport with listener	Ambient conditions	Authoritarian/control freak	Child vs elderly	
Customs	Gender	Dialogue turn	Child vs elderly	Gender	
Race		Familiarity with system	Appearance, eg spectacles, facial hair, head and eye movement	Illness/Infirmity	
		Sedentary/active		Impairment	
		Overt/covert		Vocal tract length	
		Location		Skin colour	

Figure 2: Use of the model in practice

tion. The model depicted at Figure 1 takes into account the importance of context in the analysis of affect in speech.

A recent study by Devillers et al. included context annotation as metadata to a corpus of medical emergency call centre dialogs (Devillers et al. 2005). In that study, context information is treated as either task-specific global in nature and “Call origin” and “reason for call” are considered task-specific. The model proposed in this paper does not differentiate between task-specific and global context as the difference is seen merely as temporal, i.e. pre-determined or established at “run-time”.

Other researchers have included “discourse context” such as speaker turns (Liscombe, Riccardi, and Hakkani-Tür 2005) and specific dialog acts of greeting, closing, acknowledging and disambiguation. Inclusion in a corpus of speaker turns would be useful but annotation of every specific type of dialog act would be extremely resource intensive.

A recent report from the HUMAINE project (HUMAINE) included a proposal that at least the following issues be specified:

- Agent characteristics (age, gender, race)
- Recording context (intrusiveness, formality, etc.)
- Intended audience (kin, colleagues, public)
- Overall communicative goal (to claim, to sway, to share a feeling, etc.)
- Social setting (none, passive other, interactant, group)
- Spatial focus (physical focus, imagined focus, none)
- Physical constraint (unrestricted, posture constrained, hands constrained)
- Social constraint (pressure to expressiveness, neutral, pressure to formality)

but went on to say,

“It is proposed to refine this scheme through work with the HUMAINE databases as they develop.”

Millar et al. developed a methodology for the design of audio-video data corpora of the speaking face (Millar, Wagner, and Göcke 2004) in which the need to make corpora re-usable is discussed. The methodology, although aimed at corpus design, takes into account the need for *speaker* and *speaking environment* factors. In contrast to the HUMAINE report mentioned previously, the model presented in this paper treats agent characteristics and social constraints separate to context information. This is because their effect on discourse is seen as separate topics for research. Two things are evident:

1. That context is extremely important in affect; and
2. Defining context annotation is in its infancy.

5.1.2. Agent characteristics

As Scherer points out (Scherer 2003) most studies are either speaker oriented or listener oriented with most being the former. This is significant when you consider that the emotion of someone labelling affective content in a corpus could impact the label that is ascribed to a speaker's message.

Not much attention in the literature has been given to the role that agent characteristics such as personality type play in affective presentation which is surprising when one considers the obvious difference in expression between extroverted and introverted types. Intuitively, one would expect a marked difference in signals between speakers. Again intuitively, one would think that a baseline of the person's personality type would be of great benefit in applications that monitor an individual's emotions.

At a more physical level, agent characteristics such as facial hair, whether they wear spectacles, and their head and eye movements all affect the ability to visually detect and interpret emotions.

5.1.3. Cultural

Culture-specific display rules (Cowie et al. 2005) influence the display of affect. Gender and age are established as important factors in shaping conversation style and content in many societies.

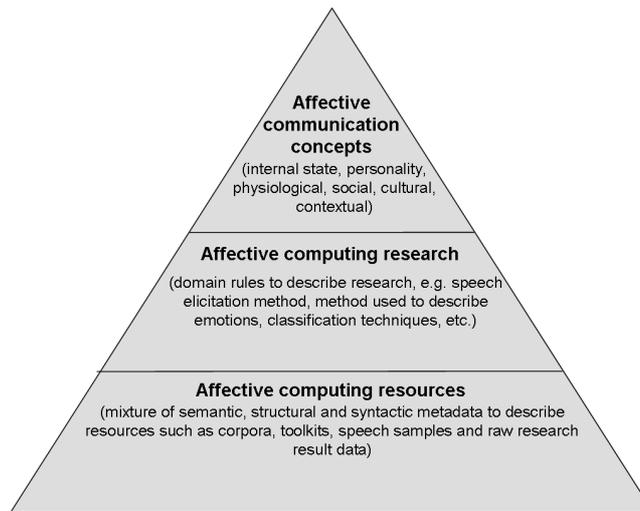


Figure 3: A set of ontologies for affective computing

Studies by (Koike, Suzuki, and Saito 1998) have shown that it is difficult to identify the emotion of a speaker from a different culture and that people will predominantly use visual information to identify emotion (Shigeno 1998). Putting it in the perspective of the proposed model, cognisance of the speaker and listener's cultural backgrounds and the context, and whether visual cues are available, obviously influences the effectiveness of affect recognition

5.1.4. Physiological

It might be stating the obvious but there are marked differences in speech signals and facial expressions between people of different age, gender and health. Millar et al. point out the habitual settings of facial and vocal organs that determine the speaker's range of image and sound possibilities. The configuration of chin, lips, nose and eyes provide the visual cues whereas the vocal tract length and internal muscle tone guide the interpretation of acoustic output (Millar, Wagner, and Göcke 2004).

5.1.5. Social

Social factors temper speech to the demands of civil discourse (Cowie et al. 2005). Affective bursts are likely to be constrained in the case of a minor relating to an adult, yet totally unconstrained in a scenario of sibling rivalry. A social setting in a library is less likely to yield loud and extroverted displays of affect than a family setting.

5.1.6. Internal state

Internal state has been included in the model for completeness. At the core of emotions is the person and their experiences. Recent events such as winning the lottery or losing a job are likely to influence emotions.

5.1.7. Examples

To help explain the differences between the factors that influence the expression of affect, Figure 2 lists some examples. The factors are divided into two groups. The first group includes those factors that are more inherent in the speaker or listener, i.e. their personality type, physiological make up and their internal state. The second is a list of

factors that modulate the speaker and/or listener's presentation of affect, i.e. cultural, social and contextual. There are many more possibilities in the second group.

5.2. The ontologies

The ontologies described in this paper are a work in progress. Figure 3 depicts the relationships between them and gives examples of each of the ontologies. Formality and rigour increase towards the apex in Figure 3. The types of users are not confined solely to researchers. There could be many types of users such as librarians, decision support systems, application developers and teachers.

5.2.1. Ontology 1 - affective communication concepts

Note that the term used is "Affective communication" and not "Affective computing". This is to enable the ontology to be expanded to other areas. The top level ontology correlates to the model previously discussed and is a formal description of the domain of affective communication. It contains internal state, personality, physiological, social, cultural, contextual. It can be linked to external ontologies in related fields such as medicine, anatomy and wordnet.

5.2.2. Ontology 2 - affective computing research

This ontology is more loosely defined and includes the concepts and semantics used to define research in the field. It has been left generic and can be further subdivided into affective computing domain at a later stage. It is used to specify the rules by which accredited research reports are catalogued. It includes metadata to describe, for example, classification techniques used, e.g. k-NN, SVM; the method of eliciting speech, e.g. acted or natural; and manner in which corpora or results have been annotated, e.g. categorical or dimensional. Creating an ontology this way introduces a common way of laying down the knowledge and facilitates intelligent searching and reuse of knowledge within the domain. For instance, an ontology just based on the models described in this paper could be used to find all research reports where:

```
SPEAKER(internalState='happy',
physiology='any',
```

agentCharacteristics='extrovert',
social='friendly',context='public',
elicitation='dimension')

Again, linking to other resources means that it does not have to be created from first principals. As an example, one resource that will be linked is the Emotion Annotation and Representation Language (EARL) which is currently under design by Schröder et al. within the HUMAINE project (Schröder 2006). EARL is an XML-based language for representing and annotating emotions in technological contexts. Using EARL, emotional speech can be described either using a set of forty-eight categories, dimensions or even appraisal theory. Examples of annotation elements include “Emotion descriptor” - which could be a category or a dimension, “Intensity” - expressed in terms of numeric values or discrete labels, “Start” and “End”.

5.2.3. Ontology 3 - affective computing resources

This ontology is more correctly a repository that contains both formal and informal rules, and data. It is a mixture of semantic, structural and syntactic metadata. It contains information about resources such as corpora, toolkits, speech samples and raw research result data.

6. Conclusions

Affective computing is important for the next generation of user interfaces and cannot be ignored. To move from being seen as something that produces things that critics only see as “cute”, it is essential that research into affective computing embraces notions such as *systematic, verifiable means* and *acquiring knowledge through empiricism*. However, the study of emotions in speech is complex and there are some difficult problems to address. It is suggested that providing a complete model of affective communication and a set of ontologies will encourage pluridisciplinary research and lay a solid foundation for future research.

Further information will be available at

www.rsise.anu.edu/gmcintyr

as the ontology is developed.

References

- Bennett, D. (2005, December). Tell lara I love her; she's attractive, charming and always there for you. *New Scientist* 188(42).
- Cornelius, R. (1996). *The science of emotion*. New Jersey: Prentice Hall.
- Cowie, R. and R. Cornelius (2003). Describing the emotional states that are expressed in speech. *Speech Communication* 40, 5–32.
- Cowie, R., E. Douglas-Cowie, and C. Cox (2005). Beyond emotion archetypes: Databases for emotion modelling using neural networks. *Neural Networks* 18, 371388.
- Cowie, R. e. a. (2000). 'feeltrace': an instrument for recording perceived emotion in real time. In *Proceedings of the ISCA Workshop on Speech and Emotion*.
- Devillers, L., L. Vidrascu, and L. Lamel (2005). Challenges in real-life emotion annotation and machine learning based detection. *Neural Networks* 18, 407422.
- Ekman, P. (1999). *The Handbook of Cognition and Emotion*, Chapter Facial Expressions, pp. 301–320. Sussex, U.K: John Wiley & Sons, Ltd.
- Ekman, P. and H. Oster (1982). *Emotion in the human face (2nd ed)* (2nd ed.). New York: Cambridge University Press.
- HUMAINE. <http://emotion-research.net/>.
- Koike, K., H. Suzuki, and H. Saito (1998). Prosodic parameters in emotional speech. *proceedings*. pp. 679–682. ICSLP.
- Liscombe, J., G. Riccardi, and D. Hakkani-Tür (2005, September). Using context to improve emotion detection in spoken dialog systems. *EUROSPEECH'05*.
- Millar, J., M. Wagner, and R. Göcke (2004, oct). Aspects of speaking-face data corpus design methodology. In *ICSLP2004*, Volume II, Jeju, Korea, pp. 1157–1160.
- Scherer, K. R. (1999). *Handbook of Cognition and Emotion*, Chapter Appraisal theory. New York: John Wiley.
- Scherer, K. R. (2003). Vocal communication of emotion: A review of research paradigms. *Speech Communication* 40, 227–256.
- Scherer, K. R. (2004). Humaine deliverable d3c: Preliminary plans for exemplars: theory <http://emotion-research.net/>.
- Schröder, M. (2006). D6e: Report on representation languages. HUMAINE.
- Schröder, M. and R. Cowie (2005). Developing a consistent view on emotion-oriented computing. Technical report, HUMAINE project.
- Shigeno, S. (1998). Cultural similarities and differences in the recognition of audio-visual speech stimuli. *ICSLP*.
- Stibbard, R. (2001). *Vocal expression of emotions in non-laboratory speech: An investigation of the Reading/Leeds Emotion in Speech Project annotation data*. Ph. D. thesis, University of Reading, UK.
- Velten, E. (1968). A laboratory task for induction of mood states. *Behaviour Research and Therapy* 6, 473–482.