EMOTION-ORIENTED COMPUTING: STATE of the ART and KEY CHALLENGES

Prof. Roddy Cowie Co-ordinator, Humaine Network of Excellence r.cowie@qub.ac.uk

1. Overview Data suggest that less than 10% of human life is completely unemotional. The rest involves emotion of some sort (http://emotion-research.net/deliverables/D3e%20final.pdf), but probably less than a quarter consists of archetypal 'fullblown emotions' (brief episodes dominated by strong feelings). Most of it is coloured by 'pervasive emotion' (feelings, expressive behaviours, evaluations, and so on, that are inherent in most human activities). 'Emotion-oriented computing' is about technology that takes account of both types. Picard's 'Affective Computing' in 1997 [1] established the field, and several EC projects made substantial contributions (NECA, ERMIS, SAFIRA, etc). The EC Network of Excellence HUMAINE now brings together many of the most active teams (and its website provides much of the background material cited here).

Emotion-oriented computing is intrinsically complex because emotional life is multifaceted. It involves affect (feelings and physical changes associated with them); cognition and conation (shifts in perception, judgment, and selection of action); interaction (registering other people's emotional status and conveying emotion to them); personality; culture; ethics; and much more. The research challenge is huge because genuinely satisfying systems need to deal with all those facets in a co-ordinated way (see section 2).

The obvious application of emotion-oriented computing is as part of the general drive to let machines interface with humans as richly as humans interface with each other. Since emotion pervades natural human interaction, emotion-oriented computing is as fundamental to the drive as speech recognition and synthesis. Its importance is likely to become clearer as allied technologies mature, and the lack of emotional 'intelligence' in otherwise humanlike systems becomes seriously anomalous. People already experience the effect to some extent in call centres that give the same bright message however long the user has been waiting.

There is no single 'killer' application for systems that support truly natural communication with humans, but it is hard to doubt that demand for them would be massive. However, the timescale is long – decades – because there are very large amounts of fundamental research to be done. It is becoming possible to identify real short term applications for the capacities we have, though. They are reviewed in section 3. Setting the right balance between exploiting limited contemporary technology and making fundamental advances is a key challenge.

2. Fundamental research: the challenges and the state of progress Until recently, the area lacked a structured research agenda. A key aim of HUMAINE was to create one. The result is that core areas have been established, and there are significant projects under way in all of them.

Descriptive resources Emotion-oriented computing needs tractable ways of describing the states that matter to it. Familiar schemes pull people to think in terms of pure fullblown emotion rather than pervasive emotion-related phenomena like friendliness, trust, distress, sincerity and mixed or time-varying emotions. Research in this area is clarifying the states most likely to matter to emotion-oriented computing (http://emotion-research.net/restricted/wp6/emolang/descry), and adapting ideas from psychology such as soft coding, dimensional representation, and appraisal theory to provide representations that are more tractable than list of irreducible categories [2]; [3],[4].

Perceiving emotion Research has explored many of the channels that people use to form impressions of each other's emotions – facial expression, paralinguistic, gesture, choice of words and actions (http://emotion-research.net/deliverables/D4b.pdf). Physiological correlates of affect also exert a special fascination [5]. High recognition rates can be obtained with acted or carefully elicited data, but the field has moved on to deal with naturalistic material [6], [7]. There it is difficult to exceed 80% success in a binary distinction (http://emotion-research.net/ws/WP4/anton-santorin.pdf). Multimodal integration seems the likeliest key to real improvement (see the special session at IEEE ICME, http://www.image.ece.ntua.gr/icme2005/).

Expressive behaviour As in perception, the existence of multiple channels is critical. Early 'Embodied Conversational Agents' (ECAs) tried to convey emotion using analyses of static faces showing fullblown emotions. The results are recognisable, but disconcerting. Research has moved on to study the rich range of signals that transmit emotion-related information in interactions, and the ways they are co-ordinated and dependent on the other party's actions. That raises topics such as eye movements, backchannelling, gesture, and 'idle movements'. Finding ways to co-ordinate these in time is a topic in its own right (http://emotion-research.net/deliverables/D6d%20final.pdf).

Emotional cognition An agent cannot engage emotionally unless it has a kind of empathy, i.e. it can understand at some level what a person's emotional state might dispose him or her to do, and how that disposition might be affected by different actions that the agent might take. Hence interfaces need to include models of central states and processes in the user that incorporate emotion. At present, several very different types of model are available – AI (using propositional representations); neurally inspired; and artificial life. Each has strengths, but it is difficult to combine them, and finding ways to do that stands out as the immediate priority (http://emotion-research.net/deliverables/D6d%20potential%20exemplars%20cog.pdf).

Complex media Emotion can be expressed and influenced not only through basic channels established by evolution, but also through music, colour, typography, and above all language. Within language there are many ways to express and influence emotion, including choice of argument, lexical selection, politeness, and humour. Work is in progress integrating all of these into the theory and practice of emotional communication (http://emotion-research.net/deliverables/D8a.pdf).

Empirical resources Progress in most areas depends on good primary records, with appropriate annotation, of people interacting emotionally with each other and machines. There is a need for both generic material (to drive fundamental research) and application-specific (to achieve tuning to particular settings). Records also need to reflect differences between people related to their gender, culture, and individual characteristics, and the context in which they are set. Techniques for both collection and annotation have been developing [3], [8], but they need to be consolidated and applied on a large scale.

Guiding system development Translating theory into product poses special problems in the area, not least because emotional aspects of response are singularly difficult to measure without changing the experience. Innovative methods are needed to gauge the kinds of innovation people may want and the way they respond to prototypes, to deliver appropriate information to designers, investors, and users, and so on. Viable products depend on combining these streams with traditional research.
Ethics Influencing people's emotions raises ethical questions, but over-reaction to the issue could stifle thoroughly desrirable developments. The field urgently needs an ethical framework that distinguishes between benign and suspect kinds of development, allied to appropriate monitoring systems. A framework based on Principlism has been proposed (http://emotion-research.net/deliverables/D10b.pdf/view?searchterm=ethical) It is hard to overstate the importance of consolidating that kind of framework and ensuring that it is generally accepted. __

3. Applying emotion-oriented systems The scope to apply emotion-oriented systems depends on the competence of the systems. By the nature of the area, that means development is gradual; but there are now areas where real applications exist, and areas where they can realistically be anticipated. This section works from the relatively established to the foresseable.
Games Emotion has become increasingly established in gaming environments, as a few examples illustrate. In **Girland** http://www.girland.com/, young girls use an emotion wheel to input their emotive states along with verbal utterances. In **The SIMS** http://thesims.ea.com/us/ , characters' actions are influenced by an emotive state, which results from their personality and events occurring during the game. In **Fahrenheit** http://www.fahrenheitgame.com/, players are able to influence the emotional state of characters indirectly. In **FinFin**, the user can influence the emotional state of a half bird, half dolphin creature directly by talking and waving. "**The Wild Divine**" http://www.wilddivine.com/ illustrates a different use of emotion: the player needs to achieve a state of relaxation or concentration to manipulate devices of the game environment.
Affective Storytelling There is well-established interest in creating stories and dramas interactively, and emotion is central to both. In the NECA project, **eShowroom** demonstrator allowed users to shape an animated car sales dialogues between virtual characters with emotions [9]. Also in NECA, **Socialite** shows simulated social interaction among emotional embodied conversational agents [10]. In the **NICE** project, users interact by speech and gesture with Hans Christian Andersen, and play out stories based on his [11]. **ActAffAct** uses BDI architecture to create emergent narratives, and it has also been extended with concepts from emotion-regulation [12]
Enhancing expression A less obvious domain, but one that users clearly value, is enriching people's ability to express emotions. Iida & Campbell [13] provided an early example with a speech synthesis system that allowed people who could not speak to use synthetic speech with emotional colouring. Fagerberg et al. [14] developed **eMoto.** a service which (literally) allowed users to add emotional colouring to text messages. The Siemens mobile phone **CX70 EMOTY** (see http://www.siemens-mobile.at/emoty) also allows users to convey their emotional states via various sensors. A strikingly simple system due to Kaye (http://alice-waters.jofish.com/writing/io-chi-short-paper.pdf) consists of a button on a PC desktop: pressing it signals to a significant other that the user is thinking of him/her. Triggering music that suits the mood is similar in some respects. For instance, the **Galaxias** system [15] generates affective music automatically in the light of the user's physiological feedback. In a different vein, systems like the **Jerk-O-Meter** allow users to monitor their own behaviour, and provide feedback to let them modify it if conveys the wrong emotions [16].
Education This has always seemed a natural application domain for emotion-oriented technologies, because emotional issues play such a large part in learning. In the **Cosmo** System [17], the agent used emotional means to achieve teaching goals. **ITSPOKE** tried to register the student's emotional states, and use them to tailor spoken tutorial dialogues appropriately [18]. On a different level, **FearNot** used emotion-driven virtual characters to teach lessons about bullying [19].
Call centres These have been used for some time as a source of data [20], and real applications are beginning to emerge. André et al. [21] provide various examples. The **T-Systems** emotion-ware voice portal [22].aims to detect user states as a way of selecting appropriate conciliation strategies and deciding whether or not to transfer the caller to a human agent. A different tack, involves using emotion detection to suggest when speech detection may be a problem in a spoken dialogue system [23].
Changing lifestyle There is developing interest in the use of emotion-sensitive agents to encourage and support changes in lifestyle. Bickmore (http://www.ccs.neu.edu/home/bickmore/JIT/) has described an emotion-oriented coach to promote exercise. Work by de Rossis on eating behaviour is reported in http://emotion-research.net/ws/wp8/proceedings-wswp8.pdf . Potential applications clearly exist in areas where non emotion-sensitive automatic systems are already used, such as helping helping elderly to exercise more. Stress relief (by providing an 'emotional gym') is also a natural target.
Therapies Computers are already providing some significant kinds of therapy. O ne of the first to incorporate emotion-sensitive techniques is Carmen's Bright IDEAS (http://www.isi.edu/isd/carte/proj_parented/description.html). It uses on drama-based interventions to help mothers of young cancer patients to develop problem solving skills. There are other areas where it is natural to consider introducing emotion-oriented techniques. For instance, LSVT therapy for Parkinson's sufferers (http://www.lsvt.org/main_site.htm) already uses a virtual agent. The UK NHS Health Technology Assessment programme has recently approved FearFighter & Beating the Blues, (http://www.cinahl.com/cexpress/hta/summ/summ622.pdf) which provide therapy for anxiety and depression, and research using VR is under way on Cognitive Bbehaviour Therapy for social phobias http://www.cs.ucl.ac.uk/research/vr/Projects/SocialPhobias/research.htm . It is natural to consider the potential of emotion-sensitive systems as support for elderly groups who need support but find standard technologies difficult.
Transport There has been interest for some time in monitoring emotion-related states in pilots and drivers [#][#] , www.image.ntua.gr/oresteia Recently Norman has added a new twist by proposing that interchange of information between car and driver could effectively take an emotion-related form http://www.jnd.org/dn.mss/theres_an_automobil.html.
Sales, marketing, and promotion The potential of emotion-related techniques in this area is obvious, but so are the ethical issues surrounding it. Both are considered in http://emotion-research.net/ws/wp8/proceedings-wswp8.pdf . Note, though, that the areas considered above do not pose the same ethical problems.

4. Forward agenda The field of emotion-oriented computing is approaching maturity. It has developed both a research agenda and a schedule or realistic applications. Europe is very strong in the area, not least because EC projects have supported individual projects and encouraged the emergence of a cohesive research community. The challenge is to sustain the momentum. Clearly research has to build engagement with emerging application areas, but in a way that nurtures fundamental research rather than diverting or sidelining it.

The natural agenda is to develop a mix of projects. Support action is needed to create large scale databases, with associated descriptive schemes, and make them genuinely accessible. It is probably possible to link development of sophisticated ECAs to major application areas, perhaps related to health and lifestyle; that would also naturally link to innovative ways of guiding system development. Fundamental research is clearly needed in two main areas – multimodal detection capabilities, and the linguistic and conceptual aspects of emotion. Interactive storytelling and drama may be suitable application areas to link to these. All of the areas need solid inputs from appropriate psychological theory, if only to prevent proliferation of ad hoc approaches to representing emotion, which would be a serious block to integration.

Not least, the ethical and privacy issues surrounding the area need to be grasped. There are some real problem areas. They need to be articulated and isolated so that the great bulk of the work, which offers genuine benefits to many people, can be allowed to develop without needless suspicion and control.

5. References

[1] Picard, R. (1997). Affective Computing. MIT Press

[2] Cowie, R., & Cornelius, R. (2003). Describing the Emotional States that are Expressed in Speech. Speech Comm, 40, 5–32.

[3] Cowie, R Douglas-Cowie, E. Cox, C.Beyond emotion archetypes: Databases for emotion modeling using neural networks, Neural Networks 18 (2005) 371–388

[4] Devillers, L., Vidrascu, L. and Lamel, L., Challenges in Real-Life Emotion Annotation and Machine Learning Based Detection, Journal of Neural Networks, vol. 18, no. 4, Jul. 2005

[5] Picard, R.W., Vyzas, E., Healey, J., Toward machine emotional intelligence: Analysis of affective physiological state, IEEE Transactions on Pattern Analysis and Machine Intelligence, 23, 10 (Oct. 2001), 1175-1191.

[6] Batliner A, Fischer K, Huber R, Spilker J & Nöth E (2003). How to find trouble in communication. Speech Communication 40, 117-143

[7] S. Ioannou, A. Raouzaiou, V. Tzouvaras, T. Maillis, K. Karpouzis, S. Kollias, Emotion recognition through facial expression analysis based on a neurofuzzy network, Neural Networks 18 (2005) 423–435

[8] Douglas-Cowie, E. Devillers, L, Martin, J-C, Cowie, R Savvidou, S Abrilian, S & Cox, C Multimodal databases of everyday emotion : facing up to complexity Proc Interspeech 2005, Lisbon 813-816

[9] Gebhard, P., Klesen, M., and Rist, T. (2004). Coloring Multi-Character Conversations through the Expression of Emotions. Proc. of the Tutorial and Research Workshop on Affective Dialogue Systems (ADS'04), Kloster Irsee, June 14-16

[10] Krenn, B., Neumayr, B. (2003). Socialite in derSpittelberg: Incorporating animated conversation into a Web-based community-building tool. In: Proc. of the 4th International Working Conf on Intelligent Virtual Agents (IVA'03), Kloster Irsee

[11] Buisine, S., Martin, J.C., and Bernsen, N.O. (2005). Children's Gesture and Speech in Conversation with 3D Characters. Proceedings of HCI International 2005, 22-27 July.

[12] Rank S., Petta P. (2005). Appraisal for a Character-based Story-World, in Panayiotopoulos T., et al.(eds.), Intelligent Virtual Agents, 5th International Working Conference, IVA 2005, Kos, Greece, September 2005, Proceedings, Springer Berlin Heidelberg, pp.495-496.

[13] Iida, A., Campbell, N., Higuchi, F., and Yasumura, M. (2003). A corpus-based speech synthesis system with emotion. In: *Speech Communication Special Issue on Speech and Emotion*, 40(1–2):161-187

[14] Fagerberg, P., Ståhl, A., and Höök, K. (2004). eMoto - Emotionally Engaging Interaction, In Journal of Personal and Ubiquitous Computing, Special Issue on Tangible Interfaces in Perspective, Springer.

[15] Kim, S. and André, E. (2004). A Generate and Sense Approach to Automated Music Composition, Proceedings of the 9th International Conference on Intelligent User Interfaces, 2004, 268--270, ACM Press, Funchal, Madeira, Portugal

[16] Madan A., Caneel R., and Pentland A. (2005). Voices of Attraction. Proc of Augmented Cognition, HCI 2005, Las Vegas.

[17] Lester J. C. et al., "Deictic and Emotive Communication in Animated Pedagogical Agents," Embodied Conversational Agents, MIT Press, Cambridge, Mass., 2000, pp. 123–154

[18] Litman, D., Forbes, K., and Silliman, S. (2003). Towards emotion prediction in spoken tutoring dialogues. In: NAACL '03 Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology, 52 – 54, Edmonton, Canada

[19] Aylett, R.S., Louchart, S., Dias, J., Paiva, A., and Vala, M. (2005). FearNot! - an experiment in emergent narrative. Proceedings, IVA 2005, LNAI, Springer

[20] Vidrascu, L. and Devillers, L, Annotation and Detection of Blended Emotions in Real Human-Human Dialogs Recorded in a Call Center, Proceedings of International Conference on Multimedia and Expo (ICME), Amsterdam, Jul. 2005

[21] E André, L Dybkjær, W. Minker & P. Heisterkamp (eds.): Affective Dialogue Systems, 2004, Proceedings Springer 2004

[22] F. Burkhardt, M. van Ballegooy, R. Englert, R. Huber (2005): An Emotion-Aware Voice Portal, ESSP.

[23] Riccardi, G. and Hakkani-Tür, D. (2005). Grounding Emotions in Human-Machine Conversational Systems. In: Mark Maybury, Oliviero Stock, Wolfgang Wahlster (eds.) Intelligent Technologies for Interactive Entertainment: First International Conference, INTETAIN 2005, Madonna di Campiglio, Italy, November 30 – December 2, 2005. Proceedings, pp. 144-154

[24] McCall, J, Mallick, S, Trivedi, M Real-time driver affect analysis and Tele-viewing system Vehicles Symposium, Proc IEEE, Jun 9-11, 2003, 372-377

[25] Hadfield, P., Marks, P., 2000. This is your captain dozing. New Scientist 1682267, 21