

humaine

D11d: Final report on WP11

Workpackage 11 Deliverable



Date: 30th January 2008

IST project contract no.	507422
Project title	HUMAINE Human-Machine Interaction Network on Emotions
Contractual date of delivery	<i>January 30, 2008</i>
Actual date of delivery	<i>February 14, 2008</i>
Deliverable number	D11d
Deliverable title	Final report on WP11
Type	Report
Number of pages	15
WP contributing to the deliverable	WP 11
Task leader	OFAI
Author(s)	Paolo Petta
EC Project Officer	Philippe Gelin

Address of lead author: Paolo Petta,
Österreichische Studiengesellschaft für Kybernetik (OFAI),
Freyung 6/6
A 1010 Vienna
Austria (EU)

Table of Contents

1	THE PLACE OF THIS REPORT WITHIN HUMAINE	5
2	WORKPACKAGE REMIT, MEMBERSHIP AND STRUCTURE	6
2.1	The HUMAINE Community	6
3	SCOPE AND STRUCTURE OF THE HUMAINE HANDBOOK	8
4	SHORT OVERVIEW OF THE HANDBOOK CONTENT	9
4.1	Theories and Models of Emotion	9
4.2	Signals to Signs	9
4.3	Databases.....	9
4.4	Emotion in Interaction.....	10
4.5	Emotion in cognition and action	10
4.6	Emotion in communication	10
4.7	Usability of emotion-oriented systems.....	11
4.8	Ethical issues	11
5	SUMMARY OF WORKFLOW AND CURRENT STATUS.....	12
6	CONCLUSION	13
7	WP11 DELIVERABLES (ALL OF THESE ARE AVAILABLE ON THE PORTAL)	15

1 The place of this report within HUMAINE

The final workpackage reports in HUMAINE have been designed as brief statements that complement other sources. On one hand, substantive content – particularly technical content – has been reported in deliverables throughout the project. On the other hand, several potentially overlapping sources are part of the final reporting process. In particular, the activity report for the final period sets out the goals achieved during the final period; and the final report for the project as a whole includes a relatively non-technical account of what the workpackage has achieved.

This core of this report is shared with the section on WP11 in the final report for the project as a whole. That consists of a short, non-technical summary of what the workpackage has achieved. However, there are details that do not belong in a report intended for the general reader (such as information about the participating institutions, meetings organised, etc). They are covered here.

Workpackage members invested a substantial amount of effort in the deliverables, and they form a continuing resource available to the community on the HUMAINE portal. They are listed for completeness. In some senses it would make sense to do the same for publications, but they are too difficult to allocate to individual workpackages. A full list is given in the final report for the project as a whole.

2 Workpackage remit, membership and structure

Over four years of intense and continuous interdisciplinary discussion and joint efforts, the work of the HUMAINE Network of Excellence has produced important contributions to laying out the foundations for the development of systems that can register, model and/or influence human emotional and emotion-related states and processes: emotion oriented systems. While there is widespread interest in such systems that may be central to future interfaces, the theory that should underpin their development is fragmented, dispersed across many disciplines, incomplete, and often completely unfamiliar to the stakeholders, who comprise the technological, end-user, policy-making communities. This theory is extensive because emotion in the broad sense pervades human communication and cognition. Human beings have positive or negative feelings about most things, people, events and symbols. These feelings strongly influence the way they attend, behave, plan, learn, and select. The feelings are conveyed e.g. in faces, voices, gestures, postures, and pace of interaction; and people judge others by the way they respond to such signals. They seek out situations that are emotionally rich (not necessarily pleasant), and are persuaded as much by emotional appeal as by factual content.

The real potential of emotion oriented systems technologies and related best practices still need to be developed, as clearly reflected in the highly heterogeneous and often dissatisfactory quality of contributions submitted to key scientific conferences such as ACII (Affective Computing and Intelligent Interaction), IVA (Intelligent Virtual Agents), and other international workshops and symposia. HUMAINE is aimed at bolstering sound long-term development by clarifying the scientific, cultural and ethical underpinnings of emotion-oriented computing, developing accessible literature, and working towards standards based on that understanding.

2.1 The HUMAINE Community

Reflecting the particularly difficult type of challenge that has been taken on, HUMAINE is a spearheading interdisciplinary exercise of intellectual integration carried out by leading experts from key disciplines. It has already succeeded in setting a global example that is now being followed up in the USA and Asia: a foundational research oriented project entitled to candid assessments of the theoretical possibilities and impossibilities, of today's technical viabilities and limitations, and of the ethical implications of development and deployment of emotion-oriented systems technologies.

By its nature, the HUMAINE project is aimed at nurturing and strengthening a living community of practitioners, at establishing the field of emotion-oriented systems as a widely accepted and well articulated domain of activity that interested parties can readily locate and relate to. Achievement of the goals of this social enterprise comprises on the one hand the forging of a community, based on the commitment of leaders in hitherto separate fields to establishing a working dialogue. This of course need not — and in fact: cannot and should not — be characterised by reaching full agreement on all positions raised. But it does encompass the development of respectful awareness— if not fuller understanding — of the differing procedures, priorities and considerations. It does also entail the at times particularly laborious construction of a shared vocabulary. And it does also — as importantly as inevitably — presuppose and require the readiness of all parties involved to alter established procedures and to jointly face and overcome the risk of “professional suicide”: the all to real threat

encountered by anyone daring to sail into waters offshore their community's consensual practices. Luckily, already very early on, at the first thematic workshop in the first year of HUMAINE hosted at the University of Geneva, the reality of both aspects was confirmed: I.e., not only the (not deliberate) provocation posed e.g. by the computer animation experts demanding for emotion psychologists to meet their every requirement in terms of complete theories amenable to ready deployment in their applications; but also the frank and patient explanation about the relation between the scope of psychological emotion theory and the full extent of emotional life¹, the ensuing discussion identifying the new territory "in-between and beyond" the reach of today's subject matters, and the decision to carry out such joint research.

Within HUMAINE, progress with community-building, the development of joint vocabulary and the identification of relevant topics of debate, is well documented with the project deliverables, and perhaps most clearly with the proceedings of the individual workshops: Originally, these were planned as events subserving a specific traditional community at a time, but already with the very first workshop just referred to, the consequences of the functioning construction of a new community became evident, and broad involvement of disciplines became the rule. These products have been feeding and have helped coordinating the activities of the community.

Sustenance of the HUMAINE enterprise beyond the lifetime of the EU-funded project is structured into different measures: regarding the human resources, the HUMAINE Association was established at the last plenary meeting in Paris in June 2007, and its executive committee elected shortly thereafter during the ACII 2007 conference in Lisbon in September. The other essential output going alongside, as a single and (just) manageable product, is the HUMAINE Handbook.

¹ It should be noted that of course there was no dearth of examples of the converse distribution of roles, as well as of the involvement of other disciplines, either.

3 Scope and Structure of the HUMAINE Handbook

The HUMAINE Handbook depicts the full scope of the Network of Excellence. It provides an authoritative and self-contained reference covering current psychological and philosophical theory and the consolidated state-of-the art in engineering of emotion oriented systems. The main goal is not to describe specific systems, but to use concrete pieces of work to motivate and explain the application of design and implementation methods and technologies. The Handbook provides a solid foundation for teachers, students, and researchers in the rapidly growing area of emotion-oriented systems. At the same time, it can serve as a comprehensive introduction and reliable source of information for a broader audience of commercial developers; end-users; and policy makers.

The Handbook comprises over thirty contributions organised in eight content sections, each with an area editorial wrapping up the content of each section; touching upon important material not covered by the contributions; and relating the content of the section to the remainder of the Handbook. Individual contributions, each section, and the Handbook as a whole are aimed at being as self-sufficient as possible in their respective domains. At the same time, the Handbook offers rich connections to suggested further reading, helping the readers in the pursuit of their particular interests. The annotated pointers to further material refer not only to individual papers, but e.g. also to publication series and other communities active in areas related to the specific topics.

Each contribution is a targeted article, addressing a topical aspect of the section area while considering the cross-section connections and related perspectives. In accord with the target readership, contributions are aimed to meet divulgation and education goals and clearly expose their connection to relevance for applications. Given the fragmented state in the field, particular care is taken to provide working definitions for terminology and to relate all use of emotion-related terms to the vocabulary as developed and provided by the theory section. Important connections across sections are identified explicitly. Where appropriate, authors include key researchers from outside the constituency of the Network of Excellence (cf. developed table of content, below).

Each section provides a commented bibliography including pointers to sources of further information, as well as contributions to the overall glossary. The glossary and a carefully edited index are essential means for accessing and exploiting the substantial content provided, where each article necessarily cuts across the borders of hitherto distinct established disciplines. All contributions undergo peer reviewing; the area editors play a pivotal role in achieving a high level of consistency and cohesion, both across the contributions of their section as well as across the whole Handbook.

4 Short Overview of the Handbook Content

The Handbook begins with a general introduction explaining the genesis, the overall scope, and the target audience of the Handbook. After that, it is divided into topic-based sections. These are outlined here.

4.1 Theories and Models of Emotion

“Theories and Models of Emotion” provides an up-to-date panorama of the landscape of emotion theories, as a grounding of the integrative effort of the HUMAINE network. The most important aspect of this endeavour is the discussion of the viabilities of the mapping of models to the actual needs of the network. Given that models cannot be employed indiscriminately for all purposes, the reader is not provided with just a simple listing of available theories, but rather the goals pursued centrally with the different modelling approaches are detailed and compared critically. This section aims at providing a theoretical point of entry to the rest of the material presented in the handbook: readers are presented the theoretical implications that underlie the applied effort described in the other sections. With respect to the integrative effort undertaken by the handbook to bridge the gap between engineering perspectives and emotion theories, the community interested mainly in the theorization of emotion — psychologists, neuroscientists, and cognitive scientists alike — is provided with a picture of the most up-to-date use of technology to study emotional processes.

4.2 Signals to Signs

The section *“From Signals to Signs”* illustrates the requirements for successful signal processing, especially “in the wild”, such as from TV footage or webcams, and the capabilities and challenges of current, dynamic affect recognition techniques. It also provides links to other sections in terms of the input that feature extraction and affect recognition can provide to synthetic environments, in particular to Embodied Conversational Agents and robots, while also explaining what kind of contextual and other background information is required for affect recognition to work, and how the scope of such functionalities is delimited.

4.3 Databases

Databases are central to the development of systems that use human-like channels of communication. As a result, they have been part of emotion-oriented computing since research in the area began. They have also been recognised as an area where there were major problems to be dealt with and an area where there is a shortage of available relevant databases. For these reasons, one of the key aims of the section on *“Data and Databases”* is to convey why research that uses databases in emotion-oriented computing needs to think about database issues rather than reflexively importing habits derived from other fields; and to identify the principles that contemporary research suggests are relevant to selecting and constructing databases in this particular area. The section also provides practical techniques for both data collection and data labelling. It finishes with a description of the exemplar HUMAINE database, which is a resource to which the community can go to see and hear the forms that emotion takes in everyday action and interaction, and to look at the tools that are relevant to describing it.

4.4 Emotion in Interaction

With evident connections to the previous parts of the HUMAINE Handbook, The section “*Emotion in Interaction*” investigates the study of emotion in the context of interaction using Embodied Conversational Agents. These computer characters are embedded with computational models of emotion used to generate expressive behaviour, in particular facial expressions; speech; and gestures, for the purposes of interacting with each other and human users. This approach is particularly suited to emotion research, where theories may often seem somewhat abstract, dry or inaccessible to those who are not experts in that domain: Driving embodied conversational agent behaviour from computational emotion models not only conveys the advantages of other computational modelling approaches in terms of validating and identifying gaps in theory: Importantly, it also makes the theories and models truly come alive by exposing the human-like aspects to which we can relate. In doing so, it empowers end-users at large to understand and judge advances in this key field for themselves, regardless of background or skill level.

4.5 Emotion in cognition and action

Complementing theoretical and implementation-oriented efforts covered in other parts of the Handbook, the next section discusses efforts aimed at improving the understanding of basic issues and open research topics regarding the involvement of emotions in “*Cognition and Action*”. Acknowledging the diversity of conceptual and computational models and frameworks used to model emotional systems, contributions cover a range of analytical perspectives and computational modelling methods: Grounded “lower-level” mechanisms as embodied in robots are assessed with respect to their potential of meeting defining characteristics of emotional functionalities; as an example of transfer from natural to synthetic environments, the key role of novelty detection in emotional processes as identified in cognitive neuroscience research is discussed along with an evaluation of the deployment of derived computational models in virtual environments; challenges “higher-level” symbolic methods have to face are illustrated with the issue of how to accommodate both logically sound and associative adaptations within a single framework; at the social level, the not dissimilar problem of how transactional affect management and affective rationality can be reconciled is examined. In a contribution reaching out to various other handbook sections, the importance of assessing the “emotion potential” of any envisioned scenario so as to gain a clear understanding of the emotional capabilities a given application asks for, as well as a viable means enabling a comparative study of different models is argued for. Practitioners will be able to gain insights on how to build different types of emotion-oriented architectures and systems, what types of tasks the different approaches might be better suited for, and their current limitations. Educated end-users will be particularly interested by the more sociological and critical reflections, both of which bring a new perspective regarding our relation to emotion-oriented artefacts, which are to become more and more pervasive in our everyday life.

4.6 Emotion in communication

The handbook section on the role of emotions in *Communication* provides a comprehensive discussion of critical issues for building persuasive — or otherwise emotion-inducing — computational systems. After a clarification of conceptual and process aspects of persuasion, including a survey of emotional and non-emotional types of persuasive strategy, the generation of verbal dialogical behaviour and approaches to computational modelling of non-

verbal behaviours such as facial expressions, gestures, and attentional feedback mechanisms are discussed, as well as their integration in multimodal interfaces, and experiences gathered with deployed persuasive embodied conversational agents on the web. As a particular phenomenon with significant persuasive potential, the state of the art in computational modelling of humour, and approaches to generation and recognition of humour are discussed in detail.

4.7 Usability of emotion-oriented systems

As explained earlier, HUMAINE directly addressed the lack of unifying theories and models of behaviours characterising the field of emotion-oriented systems by discussing and where possible integrating the contributions of researchers from a range of different disciplines. The next step that needs to be taken is to transform all these theoretical insights into viable, working affective computing applications. To do so, design and evaluation methods are required that can help steer the design process. User-centred design is one of the predominant movements within human-computer interaction (HCI) and has produced good results so far. The section of the handbook on *Usability* reviews and explains existing and novel user-centred design methods of relevance to the specific domain of emotion-oriented systems. A major contribution lies in moving beyond the functional requirements on applications, towards design qualities such as aesthetics, experiences and reflection — a movement that has been termed the “third wave of HCI”. This movement is expected to exert a major impact on the development of affective computing applications in the years to come.

4.8 Ethical issues

A particular concern with HUMAINE has been to ensure solid and practical ethical support for all activities of the Network of Excellence: to this end, an Ethical Committee was established, with the primary task of monitoring that all human research received clearance from the involved institutions, but also providing clearance for a (surprisingly large) number of institutions not yet equipped with a fitting infrastructure. At the same time, these rather unexpected circumstances provided a welcome opportunity to engage in a close and lasting dialogue with the consortium members, firmly instilling a working understanding of the importance of the support provided by ethical committees, and of the importance of getting a qualified third party assessment of the ethical implications of a proposed piece of work. In addition to a summary of the lessons learned by the ethics committee during the four years of HUMAINE, the Handbook section on Ethics and Good Practice provides an introduction into the applied ethics framework of Principlism, and its application to the specific challenges posed by the domain of emotion-oriented systems. Examples of theoretical and application-oriented analyses derived from this framework — the former discussing the notion of user autonomy, and the latter reviewing the ethical implications of persuasive systems (thus relating also to the section on Communication) — illustrate the practicality and relevance of the approach.

5 Summary of Workflow and Current Status

Work on the Handbook started with a broad discussion about the general approach to take, and it was decided that the most appropriate organisation reflecting the state of the field should mirror the HUMAINE workpackages. Appointment of area editors was the next step: most of these are close collaborators of workpackage leaders; in single cases, the very workpackage leaders decided to take over this responsibility. Accompanying the cycle of thematic workshops organised by each HUMAINE workpackage from late 2004 until very early 2006, an electronic selection of valuable background material in terms of papers presented at the workshops complemented with contributions of consortium members to other events was developed on the HUMAINE portal, and a snapshot release was submitted for reviewing as HUMAINE deliverable D11a end November 2005. This kind of “editors’ choice” helped the area editors, but also the HUMAINE constituency, to identify the main issues to be addressed in the Handbook, as detailed in the previous section of this report.

In the following, a set of more specific criteria for actual contributions to the handbook were defined. With regard to the contributions per area, these included: the number and size of contributions; overall section structure, with an editorial; targeted articles; a bibliography with other pointers to sources to further information; and contributions to an overall glossary. While main authorship was to remain within the consortium, invitation of external authors was not excluded. In authoring, a target readership not limited to researchers in emotion-oriented systems, but also including the broader community of stakeholders in emotion-oriented technology development, deployment, and use was to be taken into account. Accordingly, the style of contributions should aim to clearly expose the relevance of the content covered and to provide sufficient information (including working definitions for terminology, of particular importance given the fragmented state of the field) for the texts to be self-sufficient. Coherence of material within and across sections was emphasised as key criterion: clearly, with consistent use of vocabulary. At the same time, this obviously would not imply full agreement of perspectives on all topics addressed; but dedicated effort was to be taken to resolve disagreements identified, and otherwise an explicit clarification of the nature and causes of such divergences should be put down. All contributions would undergo peer reviewing.

Based on the material collected in D11a, these criteria, and an increased perception of the HUMAINE community as developed over participation at the different thematic workshops, the area editors fleshed out a *preliminary table of content*, mainly in a bottom-up manner — i.e., from the perspective of what the individual workpackages considered to be important contributions. The result of these activities was published as deliverable D11b in early 2006. This deliverable served as references for the following activities, comprising a cross-checking of the contributions proposed within the full editorial board, with the aim of identifying and correcting redundancies and omissions in coverage, and detecting opportunities for cross-referencing; confirming and improving the range of contributing consortium partners, to which end the whole HUMAINE constituency was invited to comment on the draft table of content and propose contributions; identifying and resolving over-commitments of individual authors; and finally also entering into contact with a range of candidate publishers identified.

In the further work on individual contributions, particular attention was dedicated to the anticipated tension between the scientific orientation of a large part of the consortium (thus aiming to report on latest achievements) and the main purpose of the Handbook as a solid reference work of more extended lifetime. Ensuring that an overall balance was preserved across the different sections; and respecting the overall practical length limitations were

further key editorial challenges in this phase. As a result, a number of proposed contributions had to be excluded, e.g., because of their overly specialised content, or because of space limitations. Many of these articles, however, were readily considered for separate projects, further enriching the set of publications delivered by the network. The resulting developed outline of Handbook contents was published as deliverable D11c end 2006.

Work on the individual contributions was pursued with emphasis in 2007, with assistance by the editorial staff. For example, given the particularly widespread immediate relevance of the handbook section covering theory to all of the other content, alignment of these contributions with the other sections of the handbook was verified with special care, early versions of the article addressing emotional vocabulary were disseminated, and explicit personal contacts were established across authorships. Specific procedures for collection of meta-information from articles — in particular, keywords and glossary entries — were defined as part of the workflow that would lead to the completed Handbook manuscript, which also includes the peer reviewing process and the final post-processing of articles. At the same time, the development of the field of emotion oriented systems research was monitored, the existence of the Handbook project was advertised, and potential conflicts with other enterprises submitted to the publishers contacted were handled.

Almost inevitably for a project of the size of the Handbook, some delays were incurred over the nominal production schedule set by the general editor. The most important causes were the need to prioritise among and rank multiple Network-related activities of the authorship, where postponing work on the Handbook was also indicated because of the resulting possibility to exploit the more durable outcomes of analyses of more scientific results of more limited lifetime (as explained) but at times also necessary because of the very success with Network activities, leading to important career steps — dissertations and professorships; some changes in composition of authorship and scope of articles, with a few authors not being able to follow up on commitments taken and some further useful contributions requiring some restructuring efforts being identified only at the time of authoring; and different degrees of efficiency (but also plain luck) in managing to constrain the overhead of coordinating the at times sizeable authoring teams (in terms of individuals, but also of different institutions). Even so, at the time of writing, 27 out of 32 handbook contributions have reached the stage of peer reviewing; the timelines of the remaining articles has been clarified, with measures in place for swift further processing of these texts as they become available. All editors are now focussing on authoring of the editorials, before taking on the final steps (mostly formal: verification of formatting, consolidating of the bibliographies and the indices of names and terms), in the production of the handbook.

6 Conclusion

While reference works do exist for individual component areas of the domain of emotion-oriented systems, such as neuroscience and psychology; multi-modal interface technologies; embodied conversational characters; or user-centred design methods for advanced human-computer interaction technologies; to date there is no comprehensive reference that integrates emotion theory, technologies, and important aspects in good practice and ethics in a coherent, authoritative, and accessible manner. At the same time, the full scope of the HUMAINE Network could not possibly have been consolidated in a single Handbook. In this regard, the HUMAINE Handbook forms a hub for the wide range of more focussed and specialised publications — books, special journal issues, conference papers, and deliverables — produced by the project's individual workpackages, but also by the coalescing community at large.

The intended audience of this advanced intelligent human-computer interface technology book consists of researchers, students, and engineering groups, as well as a more general broader audience. The book is meant to serve as unique consolidated reference for the curricula in the topical domain being established world-wide. Given that students are expected to represent an important segment of the readership, care has been taken to achieve a good compromise of quality (including colour graphics) and (financial) accessibility. We already discussed how the HUMAINE Association is set to progress the success story of the HUMAINE Network of Excellence, and the role assigned to the HUMAINE Handbook in supporting the sustained impact aimed for. The mission of the HUMAINE Association established in the final year of the HUMAINE project includes also the continuation of this particular kind of publication activity of reference material initiated with the HUMAINE Handbook. These plans include a furthering the coverage of the emotion oriented systems domain (e.g. regarding the area of Emotions and Music), and providing updates to published content.

7 WP11 Deliverables (all available on the portal)

Del. no	Deliverable name	WP no.	Delivery due	Completed
D11a	Electronic collection of papers to be developed into handbook chapters	11	21	24
D11b	Preliminary outline of handbook contents and proposals for progressing the handbook	11	24	24
D11c	Developed Outline of Handbook Contents (Confidential)	11	36	37
D11d	Final report on WP11	11	49	49