E-VOX: A Socially Enhanced Semantic ECA

Joaquín Pérez
DIIS - University of Zaragoza
María de Luna, 1
50018 Zaragoza, Spain
jpmarco@unizar.es

Eva Cerezo
DIIS - University of Zaragoza
María de Luna, 1
50018 Zaragoza, Spain
ecerezo@unizar.es

Francisco J. Serón
DIIS - University of Zaragoza
María de Luna, 1
50018 Zaragoza, Spain
fjseron@unizar.es

ABSTRACT

In this paper, we present E-VOX, an emotionally enhanced semantic ECA designed to work as a virtual assistant to search information from Wikipedia. It includes a cognitive-affective architecture that integrates an emotion model based on ALMA and the Soar cognitive architecture. This allows the ECA to take into account features needed for social interaction such as learning and emotion management. The architecture makes it possible to influence and modify the behavior of the agent depending on the feedback received from the user and other information from the environment, allowing the ECA to achieve a more realistic and believable interaction with the user. A completely functional prototype has been developed showing the feasibility of our approach.

Keywords

Embodied Conversational Agent, Cognitive Architecture, Affective Computing, Social Interaction

1. INTRODUCTION

The concepts of intelligent agent and embodied conversational agent (ECA) have emerged and progressed during the last years. In particular, an ECA is a virtual character with conversational and learning skills, realistic body expression, and the ability to perceive the environment, to reason about it and act accordingly [7]. The incorporation of these characteristics in ECAs involves a variety of challenges, including physical appearance, cognitive modeling, abilities to interact with humans and other agents, and ethical issues. Embodied agents also provide a social dimension to the interaction with users that both raises the believability and perceived trustworthiness of agents, and increases the engagement of the user with the system [24]. Presentations given by an ECA are perceived as more entertaining and less difficult [35]. Some users prefer to disclose intimate or personal experiences with an ECA than with a human [19].

There is nowadays a consensus among researchers in fields such as psychology and neuroscience about the close interaction between cognition and emotion in humans [17]. Particularly, evidence shows that emotions play an essential role in cognitive functions such as decision making, learning, and planning [8]. These research results have led to the concept of *affective computing* [27], which refers to any type of computing related to emotions or other affective phenomena. In recent years, this research area has reported numerous contributions [34, 29]. There is a consensus about considering affective aspects as essential to achieve human-like social behavior. Understanding and modelling social interactions in a way that ECAs can display human-like abilities is the goal of a promising research domain called *social signal processing* [37]. Realistic social interaction with an agent requires that the agent considers emotions of the user and itself, to enable advanced affective features such as empathy, and to combine verbal and nonverbal behavior in a seamless, unified way [1].

An approach to support all these aspects in ECAs is the use of cognitive architectures. A cognitive architecture can be defined as a scheme or pattern for structuring the functional elements that make an intelligent agent as a whole [18], so it simplifies the design of an integrated system in which all desired features can be included. Cognitive architectures seem appropriate to provide ECAs with capabilities for cognitive processing such as learning, decision making, planning, and perception. Nevertheless, in spite of the great advances achieved in the last years in the development of ECAs, there is a lack of architectures that effectively combine *cognitive* and *emotional* aspects. This type of cognitive-affective architecture must be designed to provide mechanisms that model the interaction between cognitive and affective processes and to achieve a crucial requirement in the development of ECAs: *real-feel human-computer interactions*.

On the other hand, in the last years, a huge amount of information has become available thanks to the Web. Computers are good at processing huge amounts of data, but most contents in the Web have been (and still are) mainly human-oriented, i.e. users have to interpret the meaning of the information that is exposed to them. Therefore, it was proposed to move into the Semantic Web [4], a Web where the semantics of the different resources are made explicit, thus allowing computers to process automatically that structured information. This has raised new opportunities to develop intelligent agents that help users find what they are looking for, as users are often overloaded with vast amounts of information, most of which is inaccurate, misleading or unsolicited. Among the Semantic Web applications, DBpedia is one of the most representative and active of them [5]. It provides a semantic entry point to the Wikipedia that is able to perform intelligent searches and return only results relevant to a particular knowledge domain.

In this paper, we present E-VOX, an emotionally enhanced semantic ECA designed to work as a virtual assistant, providing the user information available through DBpedia. An important aspect of this agent is that it uses extensively a cognitive-affective archi-

DAA'16, November 16 2016, Tokyo, Japan

DOI: http://dx.doi.org/10.1145/3005338.3005340

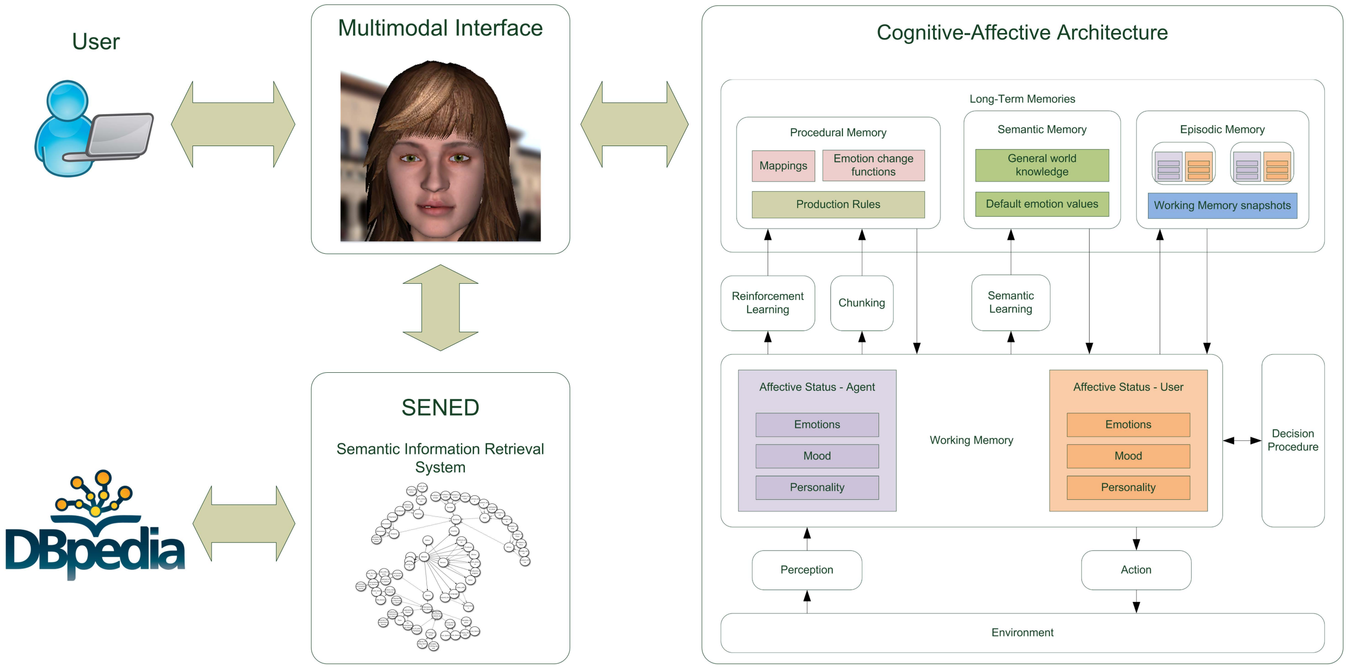

Figure 1: E-VOX: General architecture

tecture based on Soar and improved with an emotional model inspired by ALMA. This architecture allows the ECA to take into account features needed for social interaction such as learning and emotion management.

The paper is structured as follows. In Section 2, related work is presented. Section 3 is devoted to present the architecture of E-VOX system. In Section 4, the influence of the cognitive-affective architecture in the behaviour of the system is described whereas in Section 5 user-agent interaction is detailed. Finally, concluding remarks and future work are presented in Section 6.

2. RELATED WORK

In this section, we briefly present related work about cognitive architectures, affective models and ECAs.

2.1 Cognitive architectures and affective models

There are different cognitive architectures that can be used to develop intelligent agents. In [26] a review of some of them can be found. Most cognitive architectures lack some essential features (e.g., episodic memory and affective mechanisms), but nevertheless they provide a solid environment for adding this kind of requirements. ACT-R [2] and Soar [16] seem appropriate architectures to develop ECAs, given their maturity, wide range of features (e.g., different memory systems and learning mechanisms) and available free implementations in different programming languages. Another possibility is to use frameworks such as BDI (Belief, Desire, Intention), that has been used in different applications [38]. However, cognitive architectures tend to include advanced features such as reinforcement learning or specific use memories, not present in such frameworks.

On the other hand, a critical issue in the modeling of affective aspects is the lack of consensus about the definition of key concepts such as the term *emotion* [14]. Emotions can be defined as the result of the subjective interpretation of a meaningful event for a system. Several emotion models are based on the concept of *appraisal*. These models propose that agents constantly evaluate perceived events using a set of *appraisal dimensions* which are mapped to a single emotion or a range of related emotions. Among these models, OCC [25] is probably the most widely used one. However, given its complexity, it is usually simplified and only a few of the 24 emotions originally proposed are considered in most applications. Other models have also been used, such as Roseman's [31], Scherer's [33] and Cathexis [36]. These emotional models may be combined with other aspects such as mood (Soar-Emote [20]) and personality (ALMA [12]) to reflect long-term emotional states, as well as with features such as learning (FLAME [10]), empathy (FAtiMA [9]), and coping (EMA [21]) to build modern, believable, and complex intelligent agents.

2.2 ECAs

There is a large volume of literature that explores the applications and the development process of ECAs. ECAs have been used in diverse application domains, as e-Learning [21], games [3] or information assistants [28] that help the user to find useful information. This kind of agents are what some authors describe as *assistant conversational agents* (ACAs) [32].

In some of these applications, special effort has been made to work on affective issues.

The **VirtualHuman** System [30], a 3D interactive application, includes the affective model ALMA [12] to provide it with emotionally processed information useful to improve its conversational abilities.

The **Mission Rehearsal Exercise** project, a virtual-based training program intended to teach soldiers how they should act in stressful situations, includes EMA [21] to influence the decision-making of computational agents.

The **PETEEI** (*PET with Evolving Emotional Intelligence*) is an interactive emotional pet that implements FLAME to show different emotional states [10]. FLAME incorporates learning components and is implemented using fuzzy logic techniques.

Table 1: Events and appraisal tags

Event	Appraisal Tags	
	Agent	User
User provides positive explicit reward	good_event	-
User provides negative explicit reward	bad_event	-
A new information search is done	-	good_likely_future_event
The user likes the information that the agent has found	good_act_self	good_act_other
The user do not like the information that the agent has found	bad_act_self	bad_act_other
The user has spent some time without interacting with the agent	bad_event	-
The agent does not find any results for the solicited topic	bad_act_self	bad_act_other
The agent finds an appropriate result and provides it to the user	good_act_self	good_act_other
The agent cannot find any new information for the solicited topic	bad_act_self	bad_act_other
The agent decides to ignore a search request by the user, because it is angry	good_act_self	bad_act_other

Yuppy is a physical emotional pet capable of displaying certain emotional behaviors according to the particular situation in which it is involved. The affective model Cathexis [36] was included on its architecture.

3. E-VOX

E-VOX is an emotionally enhanced semantic ECA that helps users to find useful information. Figure 1 shows the architecture of the system. In the next subsections, the different modules of the system are explained.

3.1 Cognitive-Affective Architecture

The core of E-VOX is a cognitive-affective architecture based on Soar [16] and extended with an affective model inspired by ALMA [12].

Soar is a cognitive architecture designed to model the intelligent behavior of an agent. Soar adopts the idea that knowledge, planning, reaction to events, and learning can be integrated in a simple and homogeneous architecture. Soar includes a wide set of mechanisms such as semantic memory, episodic memory, reinforcement learning, and spatial and visual information system.

ALMA is an emotional model based on three layers: emotion, mood, and personality. Emotions represent an affective response in the short-term and tend to decay quickly once its cause is removed. Their intensity depends on the particular emotion elicited, the current mood, and personality [15]. In ALMA, the emotion generation process is based on the OCC model. Mood is defined as a durable emotional state that influences actions from an entity, representing the medium-term. Mood is expressed with three independent traits according to the PAD *temperament* model from [23]. Personality is based on the OCEAN model [22] and defined as the set of mental characteristics of an entity that makes itself unique, representing the long-term.

The combination of Soar and ALMA allows E-VOX to include and take advantage of features such as reinforcement learning, episodic memory, and emotion management. In fact, emotions are not only managed, as they influence all the behavior of the system. One critical point in every emotional management system is the set of emotions to be considered. In this case, and due to the application area chosen, the selected emotions are partially based on a previous study [11] that determined that the most frequently emotions experienced by the users during the information search on Internet are *surprise, hope, joy, distress, liking* and *dislike*. On the other hand, after the search process, users reported mainly emotions as *joy, relief, pride, regret, frustration, disgust* and *anger*. Based on this study, and adding some other emotions such as *boredom* and *fear*, we have selected the following emotions for the agent (*boredom, anger, fear, remorse, pride, hate* and *relief*) and the user (*joy, anger, hope, love/like* and *hate*). The consideration of two sets of emotions (user and agent) managed independently allows the architecture to manage the emotions and emotional reactions of the agent, while being able to simulate in a basic way the emotional state of the user, and opens the possibility of considering characteristics such as empathy.

3.2 Multimodal Interface

This module manages a multimodal interface based on a 3D ECA. It is responsible for capturing user input (currently, keyboard & mouse only, although ASR and facial expression recognition are planned for future work) and sending the information to the cognitive-affective architecture. Moreover, this module is also in charge of generating the animations of the 3D virtual agents. To enhance its communication capabilities, the agent performs different movements depending on which state of the interaction it is:

- When it is resting (there is no user interaction), the agent makes unconscious movements, which affects its head, eyes, eyebrows and body (secondary animations).
- In other interaction states, the agent makes facial gestures related to questions, emotions or answers, to emphasize the visual message that the user receives along with the bare communication information.

Besides, the facial expression and the voice of the agent (speed, volume and pitch) change depending on the current emotional state of it. Regarding the voice interaction, Loquendo TTS has been integrated as a speech synthesizer. An analyzer of the X-SAMPA phonetic symbols converts them into information easily recognizable in order to perform a correct viseme. The lexical and syntactic analyzer uses JFlex and CUP.

3.3 SENED module

This module provides access to the semantic knowledge stored in DBpedia. It processes the semantic keyword based queries that are received from the multimodal interface during the interaction with the user. It comprises the following components:

- The Lucene repository that ranks retrieved results and stores them in a cache.
- The Query Engine. It is in charge of building the queries and posing them to the data repositories, using DL reasoners.
- A Virtuoso repository, which provides access to the stored data via SPARQL.

More details can be found in [6].

4. INFLUENCE OF THE ARCHITECTURE

Adding a cognitive-affective architecture allows the agent to manage the behavior of the agent integrating every kind of information from different sources (short term memory, episodic memory, semantic memory, preference system) and incorporating the affective component in the decision procedure. For each event that needs to be treated by the system, every possible action is proposed (stored as production rules), and depending of the available information, one action is chosen, which represent the reaction of the system to that event. The details of this chosen reaction are sent to the multimodal interface.

As mentioned in section 3.1, the system manages the emotional state of the agent and the user. In the case of the agent, its personality is defined and bounded to its initial emotional state. In the case of the user, the first time that he or she interacts with the system, a brief personality test is conducted, and at the beginning of each user session, the system asks about his or her current mood, selecting one of the eight octants that the PAD model defines [23]

From there, the system manages the emotional impact of the events during the interaction (user inactivity, too many search results, asking for a new search, no search results...) for the agent and for the user. For this, each system event is classified, marking them with *appraisal tags*, emotional tags that represent the appraisal of the system regarding each significative event that it must process, following the ALMA model [12]. Some examples of events of the system are shown in table 1 along their corresponding *appraisal tag*, from the point of view of the user and the agent. The relationship between *appraisal tags* and specific emotions is not an one-to-one relationship, as an agent could react emotionally to the same event in different ways depending on its previous actions (and consequences) and its current emotional state. The appraisal tags allow a subset of the production rules of the architecture to fire in parallel, and the chosen one acts as the emotional reaction of the system. Ultimately, the processing of this reaction is what updates the active emotion list, and the ALMA values are recalculated. The user can provide feedback to the system at any time, about its own emotional state and the behavior of the agent. This allows the system to adjust the interaction and responses of the agent to fit more adequately to the expectations of the user.

5. INTERACTION WITH THE SYSTEM

As mentioned in previous section, the first time that the user interacts with the system, a brief personality test is conducted. The concept of user session has been added to the system, as a simple user account system is provided. After logging in, the agent is able to remember his or her search history, personality and last emotional state when the previous session ended.

Figure 2 shows a screenshot of the application interface. It shows the main window with the animated 3D ECA, areas for information display from agent and user, and additional controls to provide explicit feedback to the system about how satisfied is the user with the current behavior of the agent. This allows the agent to modify its behavior if it seems inappropriate, or reinforce it if it seems correct.

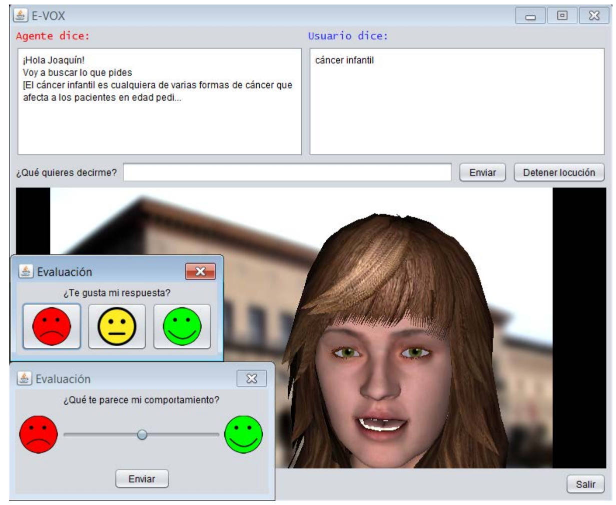

Figure 2: E-VOX user interface and feedback controls

Figure 3 shows a system window, not normally shown to the user, with controls for modifying the personality of the agent, as well as the currently active emotions, personality and current mood of the user and the agent. There are also hidden controls to turn on and off some settings of the system.

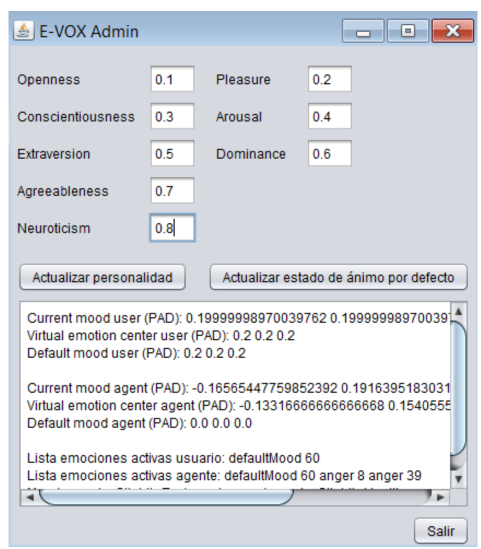

Figure 3: E-VOX control window

Some examples of interaction are shown below:

- Search attempt: The user requests the agent to search some information for him or her (*good likely future event*). Depending of its emotional state, the agent will choose to perform the task or refuse it, ignoring the user.

- Inactivity: The user stops interacting with the system for a reasonably long time (*bad event*). The agent may get angry (*Come back immediately!*), get bored (*Are you here? I'm bored...*) or worry about it (*Are you OK? Has something awful happened to you?*) depending on the current emotional state and previous interactions with the user.

- No results: The agent cannot find any results for the provided search terms from the user (*bad act self*). It may feel very sorry about its own inability, or blame the user for providing bad input, depending on the current emotional state of the agent.

- More than one result available: The ECA provides the best result available (*good act self*). If the user repeats the same request, the ECA will not provide the same result as the previous one, instead it will return the second best, third best, etc.

A first evaluation of the system is being conducted. A total of 10 persons (5 male, 5 female), aged between 17 and 58 years, have participated so far in the evaluation sessions. In each testing session, it has been carried out a pre-test with questions about age, sex, education and experience with ECAs, and a short OCEAN personality test[1] that allows to adjust the emotional model of the user and adapt subsequent interaction with the agent. After the personality test, the user performs two searches of information: a predefined search (*childhood cancer*) and a free search for any topic chosen by the user. Finally, the user fills the post-test questionnaire presented in the table 2 along with the mean values obtained. For each question, the user chooses in a five-point scale between 1 (disagree completely) and 5 (agree completely).

Table 2: Post-test questionnaire and obtained results

Question	Mean
The answers have been quick	3.4
The answers have been correct and relevant	3.4
The system has been easy to use	4.5
It has not been necessary to change how to express myself	4.4
The system has been useful	3.3
It has been entertaining to work with the system	3.5
I would like to use the system at home	2.6
I feel that the system has potential in the future	4.4

Although the results showed here are initial and correspond to a fairly small number of participants, the mean values obtained show the good general acceptance of the ECA and the future possibilities of improvement.

6. CONCLUSIONS AND FUTURE WORK

In this paper, we have presented E-VOX, an emotionally enhanced semantic ECA working as an assistant to provide useful information over Wikipedia. It is based on a powerful cognitive-affective architecture (Soar) and a complete emotional model (ALMA) that allows it to support real-feel human-computer interaction. Although this ECA is still a prototype, it is a completely functional one that shows the feasibility of our approach. Moreover, the architecture of the system can be reused to create ECAs for different scenarios.

However, we are aware that there is still plenty of work to be done. We are now considering to extend E-VOX with a number of features:

- To take advantage of other built-in mechanisms that the Soar cognitive architecture provides. We plan to use more features of the semantic and episodic memory.

- To exploit Linked Data features to provide more accurate and interesting information to the users.

- To incorporate emotional-sensing capabilities like real-time facial recognition techniques [13].

- To explore the utility of the cognitive-affective architecture in other applications.

[1]Short Personality Quiz, http://psychcentral.com/quizzes/personality.htm/

Acknowledgments

This work was supported by the Spanish "Dirección General de Investigación", contract number TIN2011-24660 and TIN2015-67149-C3-1R.

We want to thank Carlos Bobed and Eduardo Mena for the design and development of the SENED module.

References

[1] A. Aly and A. Tapus. A model for synthesizing a combined verbal and nonverbal behavior based on personality traits in human-robot interaction. In *Proceedings of the 8th ACM/IEEE International Conference on Human-robot Interaction*, HRI '13, pages 325–332, Piscataway, NJ, USA, 2013. IEEE Press.

[2] J. R. Anderson, D. Bothell, M. D. Byrne, S. Douglass, C. Lebiere, and Y. Qin. An integrated theory of the mind. *PSYCHOLOGICAL REVIEW*, 111:1036–1060, 2004.

[3] C. Becker-Asano, E. Meneses, N. Riesterer, J. Hué, C. Dornhege, and B. Nebel. The hybrid agent marco: A multimodal autonomous robotic chess opponent. In *Proc. of the 2nd Intl. Conf. on Human-Agent Interaction*, 2014.

[4] T. Berners-Lee, J. Hendler, and O. Lassila. The Semantic Web. *Scientific American*, 284(5):34–43, 2001.

[5] C. Bizer, J. Lehmann, G. Kobilarov, S. Auer, C. Becker, R. Cyganiak, and S. Hellmann. DBpedia - A crystallization point for the Web of Data. *Web Semantics: Science, Services and Agents on the World Wide Web*, 7(3):154 – 165, 2009.

[6] C. Bobed, G. Esteban, and E. Mena. Enabling keyword search on Linked Data repositories: An ontology-based approach. *International Journal of Knowledge-based and Intelligent Engineering Systems*, 17(1):67–77, 2013.

[7] J. Cassell and T. Bickmore. External manifestations of trustworthiness in the interface. *Commun. ACM*, 43(12):50–56, Dec. 2000.

[8] A. R. Damasio. *Descartes' error: Emotion, Reason, and the Human Brain*. Putnam Grosset Books, New York, 1 edition, 1994.

[9] J. Dias, S. Mascarenhas, and A. Paiva. Fatima modular: Towards an agent architecture with a generic appraisal framework. In *Emotion Modeling*, pages 44–56. Springer, 2014.

[10] M. S. El-Nasr, J. Yen, and T. R. Ioerger. Flame: Fuzzy logic adaptive model of emotions. *Autonomous Agents and Multi-Agent Systems*, 3(3):219–257, Sept. 2000.

[11] C. Flavián-Blanco, R. Gurrea-Sarasa, and C. Orús-Sanclemente. Analyzing the emotional outcomes of the online search behavior with search engines. *Comput. Hum. Behav.*, 27(1):540–551, Jan. 2011.

[12] P. Gebhard. A layered model of affect. In *4th International Joint Conference of Autonomous Agents & Multi-Agent Systems (AAMAS'05)*, pages 29–36. ACM Press, 2005.

[13] I. Hupont, S. Baldassarri, and E. Cerezo. Facial emotional classification: from a discrete perspective to a continuous emotional space. *Pattern Analysis and Applications*, 16(1):41–54, 2013.

[14] C. E. Izard. The Many Meanings/Aspects of Emotion: Definitions, Functions, Activation, and Regulation. *Emotion Review*, 2(4):363–370, Oct. 2010.

[15] M. Kipp, T. Dackweiler, and P. Gebhard. Designing emotions. *KI - Künstliche Intelligenz*, 25(3), 2011.

[16] J. E. Laird. *The Soar Cognitive Architecture*. The MIT Press, 2012.

[17] R. D. Lane, L. Nadel, J. J. B. Allen, and A. W. Kaszniak. The study of emotion from the perspective of cognitive neuroscience. In R. D. Lane and L. Nadel, editors, *Cognitive Neuroscience of Emotion*. Oxford University Press, New York, 2000.

[18] P. Langley, J. E. Laird, and S. Rogers. Cognitive architectures: Research issues and challenges. *Cognitive Systems Research*, 10(2):141 – 160, 2009.

[19] G. M. Lucas, J. Gratch, A. King, and L.-P. Morency. It's only a computer: Virtual humans increase willingness to disclose. *Computers in Human Behavior*, 37:94–100, Aug. 2014.

[20] R. P. Marinier and J. E. Laird. Computational Modeling of Mood and Feeling from Emotion. In *Proceedings of the 29th Annual Conference of the Cognitive Science Society. Cognitive Science Society*, pages 461–466, 2007.

[21] S. C. Marsella and J. Gratch. EMA: A process model of appraisal dynamics. *Journal of Cognitive Systems Research*, 10(1):70–90, Mar. 2009.

[22] R. R. McCrae and O. P. John. An introduction to the five-factor model and its applications. *Journal of Personality*, 60(2):175–215, 1992.

[23] A. Mehrabian. Pleasure-arousal-dominance: A general framework for describing and measuring individual differences in temperament. *Current Psychology*, 14(4):261–292, 1996.

[24] S. V. Mulken and E. André. The persona effect: How substantial is it? In *People and Computers XIII: Proceedings of HCI'98, Sheffield (UK)*, pages 53–66. Springer, September 1998.

[25] A. Ortony, A. Collins, and G. L. Clore. *The cognitive structure of emotions / Andrew Ortony, Gerald L. Clore, Allan Collins*. Cambridge University Press Cambridge [England] ; New York, pbk. ed. edition, 1988.

[26] J. Pérez, E. Cerezo, F. J. Serón, and L.-F. Rodríguez. A Cognitive-Affective Architecture for ECAs. *Biologically Inspired Cognitive Architectures Journal*, 2016. (to appear).

[27] R. W. Picard. *Affective computing*. MIT Press, Cambridge, MA, USA, 1997.

[28] I. Poggi, C. Pelachaud, F. de Rosis, V. Carofiglio, and B. De Carolis. *Greta. A Believable Embodied Conversational Agent*, pages 3–25. Springer Netherlands, Dordrecht, 2005.

[29] R. Reisenzein, E. Hudlicka, M. Dastani, J. Gratch, K. Hindriks, E. Lorini, and J.-J. C. Meyer. Computational modeling of emotion: Toward improving the inter-and intradisciplinary exchange. *Affective Computing, IEEE Transactions on*, 4(3):246–266, 2013.

[30] N. Reithinger, P. Gebhard, M. Löckelt, A. Ndiaye, N. Pfleger, and M. Klesen. Virtualhuman: dialogic and affective interaction with virtual characters. In *ICMI '06: Proceedings of the 8th international conference on Multimodal interfaces*, pages 51–58, New York, NY, USA, 2006.

[31] I. J. Roseman. Appraisal determinants of emotions: Constructing a more accurate and comprehensive theory. *Cognition & Emotion*, 10(3):241–278, 1996.

[32] J.-P. Sansonnet, D. Leray, and J.-C. Martin. *Architecture of a Framework for Generic Assisting Conversational Agents*, pages 145–156. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.

[33] K. R. Scherer. *Appraisal considered as a process of multilevel sequential checking.*, pages 92–120. New York and Oxford: Oxford University Press, 2001.

[34] K. R. Scherer, T. Bänziger, and E. Roesch. *A Blueprint for Affective Computing: A sourcebook and manual*. Oxford University Press, 2010.

[35] A. Serenko. A model of user adoption of interface agents for email notification. *Interacting with Computers*, 20(4-5):461–472, 2008.

[36] J. D. Velásquez. Modeling emotions and other motivations in synthetic agents. In *Proceedings of the fourteenth national conference on artificial intelligence and ninth conference on Innovative applications of artificial intelligence*, pages 10–15, Providence, Rhode Island, 1997. AAAI Press.

[37] A. Vinciarelli, M. Pantic, D. Heylen, C. Pelachaud, I. Poggi, F. D'Errico, and M. Schroeder. Bridging the gap between social animal and unsocial machine: A survey of social signal processing. *IEEE Transactions on Affective Computing*, 2012.

[38] W. Wong, L. Cavedon, J. Thangarajah, and L. Padgham. Flexible conversation management using a bdi agent approach. In Y. Nakano, M. Neff, A. Paiva, and M. Walker, editors, *Intelligent Virtual Agents*, volume 7502 of *Lecture Notes in Computer Science*, pages 464–470. Springer Berlin Heidelberg, 2012.